

STATE SPACE COLLAPSE AND DIFFUSION APPROXIMATION FOR A NETWORK OPERATING UNDER A FAIR BANDWIDTH SHARING POLICY

BY W. N. KANG, F. P. KELLY*, N. H. LEE AND R. J. WILLIAMS†

*Carnegie Mellon University, University of Cambridge and University of
California, San Diego*

We consider a connection-level model of Internet congestion control, introduced by Massoulié and Roberts [33], that represents the randomly varying number of flows present in a network. Here bandwidth is shared fairly amongst elastic document transfers according to a weighted α -fair bandwidth sharing policy introduced by Mo and Walrand [35] ($\alpha \in (0, \infty)$). Assuming Poisson arrivals and exponentially distributed document sizes, we focus on the heavy traffic regime in which the average load placed on each resource is approximately equal to its capacity. A fluid model (or functional law of large numbers approximation) for this stochastic model was derived and analyzed in a prior work [26] by two of the authors. Here we use the long time behavior of the solutions of this fluid model established in [26] to derive a property called multiplicative state space collapse, which loosely speaking shows that in diffusion scale the flow count process for the stochastic model can be approximately recovered as a continuous lifting of the workload process.

Under weighted proportional fair sharing of bandwidth ($\alpha = 1$) and a mild local traffic condition, we show how multiplicative state space collapse can be combined with uniqueness in law and an invariance principle for the diffusion [10, 20], to establish a diffusion approximation for the workload process and hence to yield an approximation for the flow count process. In this case, the workload diffusion behaves like Brownian motion in the interior of a polyhedral cone and is confined to the cone by reflection at the boundary, where the direction of reflection is constant on any given boundary face. When all of the weights are equal (proportional fair sharing), this diffusion has a product form invariant measure. If the latter is integrable, it yields the unique stationary distribution for the diffusion which has a strikingly simple interpretation in terms of independent dual random variables, one for each of the resources of the network. We are able to extend this product form result to the case where document sizes are distributed as finite mixtures of exponentials, and to models that include multi-path routing. We indicate some difficulties related to extending the diffusion approximation result to values of $\alpha \neq 1$.

*Research of F. P. Kelly is supported in part by EPSRC Grant GR/586266/01.

†Research of R. J. Williams is supported in part by NSF Grant DMS-0604537.

AMS 2000 subject classifications: Primary 60K30, 90B15

Keywords and phrases: Bandwidth sharing, α -fair, flow level Internet congestion control, proportional fair sharing, multi-path routing, simultaneous resource possession, fluid model, state space collapse, invariant manifold, workload, Brownian model, diffusion approximation, reflected Brownian motion, product form stationary distribution.

We illustrate our approximation results for a few simple networks. In particular, for a two resource linear network, the diffusion lives in a wedge that is a strict subset of the positive quadrant. This geometrically illustrates the entrainment of resources, whereby congestion at one resource may prevent another resource from working at full capacity. For a four resource network with multi-path routing, the product form result under proportional fair sharing is expressed in terms of independent dual random variables, one for each of a set of generalized cut constraints.

Contents.	2
1 INTRODUCTION	3
1.1 Overview	4
1.2 Notation and Terminology	6
2 FLOW-LEVEL MODEL	7
2.1 Network Structure	7
2.2 Stochastic Primitives	8
2.3 Bandwidth Sharing Policy	8
2.4 Stochastic Process Description	9
3 SEQUENCE OF SYSTEMS AND SCALING	10
4 FLUID MODEL	11
4.1 Fluid Model Solution	12
4.2 Invariant Manifold	12
4.3 Asymptotic Properties of Fluid Model Solutions	14
5 MAIN RESULTS	16
5.1 Multiplicative State Space Collapse	16
5.2 Conjectured Diffusion Approximation	17
5.3 Diffusion Approximation for Weighted Proportional Fair Sharing ($\alpha = 1$)	23
5.4 Product Form Stationary Distribution for Proportional Fair Sharing	26
5.5 Multi-Path Routing	33
5.6 Discussion of Conjecture for $\alpha \neq 1$	36
6 PROOF OF MULTIPLICATIVE STATE SPACE COLLAPSE	42
6.1 Crucial Estimates for Fluid Scaled Processes	42
6.2 Proof of Multiplicative State Space Collapse	46
7 PROOF OF DIFFUSION APPROXIMATION WHEN $\alpha = 1$	51
8 CONCLUSION	63
References	63

1. INTRODUCTION. We consider a connection-level model of Internet congestion control introduced and studied by Massoulié and Roberts [33]. This stochastic model represents the randomly varying number of flows present in a network where bandwidth is dynamically shared between flows that correspond to continuous transfers of individual elastic documents. This model, which we shall refer to as a flow level model, assumes a “separation of time scales” such that the time scale of the flow dynamics (i.e., of document arrivals and departures) is much longer than the time scale of the packet level dynamics on which rate control schemes such as TCP converge to equilibrium. We consider the flow level model operating under a family of bandwidth sharing policies introduced by Mo and Walrand [35] called weighted α -fair policies. Here α is a parameter lying in $(0, \infty)$. The case $\alpha = 1$ is called weighted proportional fair sharing and the case $\alpha \rightarrow \infty$ corresponds to what is called weighted max-min fair.

Assuming Poisson arrivals and exponentially distributed document sizes, de Veciana, Lee and Konstantopoulos [11] and Bonald and Massoulié [3] studied stability of the flow level model operating under weighted α -fair bandwidth sharing policies (including limiting values of α). Lyapunov functions constructed in [11] for weighted max-min fair and proportionally fair policies, and in [3] for weighted α -fair policies ($\alpha \in (0, \infty)$), imply positive recurrence of the Markov chain associated with the model when the average load on each resource is less than its capacity. Lin, Shroff and Srikant [30, 31, 39] have recently given sufficient conditions for stability of a Markov model under a back-pressure algorithm when the assumption of time scale separation is relaxed. For more general document size distributions, there are a few results for specific values of α or specific distributions or topologies that provide sufficient conditions for stability of the flow level model operating under bandwidth sharing policies [6, 8, 28, 32]. A summary of these results is provided in the introduction to [14]. In general, it remains an open question whether, with renewal arrivals and arbitrarily (rather than exponentially) distributed document sizes, the flow level model is stable under an α -fair bandwidth sharing policy $\alpha \in (0, \infty)$ when the nominal load placed on each resource is less than its capacity (see [14, 15] for some first steps in this direction). Here we restrict to the case of Poisson arrivals and exponential document sizes, for which stability is well understood.

We are interested in using diffusion approximations to explore the performance of the flow level model operating under a weighted α -fair bandwidth sharing policy when the average load placed on each resource is approximately equal to its capacity, i.e., the system is heavily loaded. We are particularly interested in manifestations of the phenomenon of entrainment, whereby congestion at some resources may prevent other resources from working at their full capacity.

There are several motivations for our work. One source of motivation lies in

fixed point approximations of network performance for TCP networks (cf. [7, 13, 36]). These approximations require, as input, information on the joint distribution of the numbers of flows present on different routes, where dependencies between these numbers may be induced by the bandwidth sharing mechanism. Similarly, an understanding of such joint distributions seems important if the performance models for a single bottleneck described by Ben Fredj *et al.* [1] are to be generalized to a network.

Another motivation is that the flow level model typically involves the simultaneous use of several resources. Due to the exponential document sizes, this model can be equated (in distribution) with a stochastic processing network (SPN) as introduced by Harrison [16]. Open multiclass queueing networks operating under head-of-the-line (HL) service disciplines are a special case of SPNs without simultaneous resource possession. For certain queueing networks of this type, it has been shown [5, 42] that suitable asymptotic behavior of critical fluid models implies a property called multiplicative state space collapse, which in turn validates the use of Brownian model approximations for these networks in heavy traffic. For more general SPNs, investigation of the behavior of critical fluid models, of a related notion of multiplicative state space collapse, and of the implications for diffusion approximations, are in the early stages of development. The analysis in this paper can be viewed as a contribution to such an investigation for models involving simultaneous resource possession. For another contribution, see the paper of Ye and Yao [44] who consider a stochastic processing network with simultaneous resource possession; in contrast to the fully heavily loaded, multiple bottleneck situation considered here, Ye and Yao consider the situation of a single heavily loaded bottleneck. A further recent contribution is the important paper of Shah and Wischik [38] who have proved multiplicative state space collapse for a class of “switched” networks with multiple bottlenecks operating under a family of scheduling policies related to the maximum weight algorithm introduced by Tassiulas and Ephremides [40].

Finally, although we restrict to exponential document sizes in this paper, we would like to relax that assumption in future work. Although this involves a significantly more elaborate stochastic model (cf. [14]) to keep track of residual document sizes (because of the processor sharing nature of the bandwidth sharing policy), knowing the results for exponential document sizes is likely to be useful for such work.

1.1. *Overview.* In this paper, we consider the flow level model with Poisson arrivals and exponentially distributed document sizes operating under a weighted α -fair bandwidth sharing policy for $\alpha \in (0, \infty)$. We focus on the heavy traffic regime in which the average load placed on each resource is approximately equal to its

capacity. We recall the definition of a critical fluid model from the prior work [26] of two of the authors; this model is a formal functional law of large numbers approximation to the flow level model. The asymptotic behavior of this critical fluid model was studied in [26]. Here we show how this behavior can be used to prove a property called *multiplicative state space collapse*. Loosely speaking this says that in diffusion scale the flow count process can be approximately recovered by a continuous lifting of the lower dimensional workload process. Given the asymptotic behavior of the critical fluid model, our proof of multiplicative state space collapse follows a general line of argument pioneered by Bramson in [5], where open multiclass queueing networks operating under certain head-of-the-line (HL) service disciplines are treated. There are some differences in setup and proof details between our treatment and Bramson's [5]. These are described in detail at the beginning of Section 6. However, we wish to emphasize that our main line of argument follows that of Bramson [5]. It is interesting to note that in contrast to prior results on state space collapse for open multiclass queueing networks, our lifting map can be nonlinear (for $\alpha \neq 1$).

The multiplicative state space collapse result leads to a natural conjecture for a diffusion approximation to the workload process. For the case of weighted proportional fair sharing of bandwidth ($\alpha = 1$), we combine multiplicative state space collapse with uniqueness in law for the diffusion [10] and an *invariance principle* [20] for semimartingale reflecting Brownian motions living in domains with piecewise smooth boundaries to obtain a diffusion approximation for the flow count process under a mild local traffic condition. This diffusion lives in a polyhedral cone. It behaves like Brownian motion in the interior of the cone and is confined to the cone by reflection (or regulation) at the boundary where the direction of reflection is constant on any given boundary face. We illustrate this diffusion approximation result for a simple two resource linear network. Then the diffusion lives in a wedge that is a strict subset of the positive quadrant. This geometrically illustrates the entrainment of resources, whereby congestion at one resource may prevent another resource from working at full capacity. We also observe how the wedge can vary with the weights. Ongoing work is directed towards establishing diffusion approximations for the workload process when $\alpha \neq 1$. We mention some of the difficulties associated with this. These center around the fact that when $\alpha \neq 1$ and the workload dimension is three or higher, although the state space for the proposed diffusion approximation for the workload process is a cone, it is not a polyhedral cone. Indeed, the cone has curved boundary faces that intersect non-smoothly and can even meet in cusp-like singularities. The current lack of a general existence and uniqueness theory (and an associated invariance principle) for reflecting Brownian motions in such domains is a major obstacle to proving the conjecture for $\alpha \neq 1$.

In the case of proportional fair sharing, i.e., when $\alpha = 1$ and all of the weights for

the bandwidth sharing policy are equal, we show that the approximating diffusion has a product form invariant measure. When the latter is integrable over the state space, our results suggest a strikingly simple approximation for the joint stationary distribution of the number of flows present on different routes under proportional fair sharing and the mild local traffic condition. In this, each of the resources of the network has associated with it a dual random variable; these dual variables are independent and exponentially distributed; and the formal approximation to the number of flows on a route is proportional to the sum of the dual variables along the route.

We also indicate an extension of the product form result to the situation where document sizes are finite mixtures of exponential distributions. Under this extension, the formal approximation for the joint stationary distribution for the number of flows present on different routes is insensitive: that is, the approximation does not depend on the distributions of document sizes other than through the means of these distributions, provided the distributions are finite mixtures of exponentials. This result complements the known result [3, 4, 33] that, for proportional fair sharing and a small class of topologies and parameters, the stationary distribution for the number of flows present on different routes is exactly insensitive: that is, the stationary distribution does not depend on the distributions of document sizes other than through the means of these distributions.

Finally, we indicate a relation to more general models with routing. There is considerable interest in multi-path routing in the Internet, and rate control schemes generalizing TCP have been proposed [19, 25]. It is known that the stability region for the flow level model may be strictly increased if multi-path routing is allowed [19, 27]. We show that our results on diffusion approximations under proportional fair sharing extend to the multi-path case. The local traffic condition becomes more difficult to verify in this setting, but if it is satisfied then our results suggest a simple approximation for the stationary distribution of the numbers of source-destination flows in terms of independently distributed dual variables, one for each generalized cut constraint.

A summary of some of the results of this paper (without proofs) was given in the conference proceedings papers [22] and [23].

1.2. Notation and Terminology. For each positive integer $d \geq 1$, \mathbb{R}^d will denote d -dimensional Euclidean space and the positive orthant in this space will be denoted by $\mathbb{R}_+^d = \{x \in \mathbb{R}^d : x_i \geq 0 \text{ for } i = 1, \dots, d\}$. When $d = 1$, we sometimes write \mathbb{R} instead of \mathbb{R}^1 and \mathbb{R}_+ instead of \mathbb{R}_+^1 . The Euclidean norm of $x \in \mathbb{R}^d$ will be denoted by $|x|$. Vectors in \mathbb{R}^d will be assumed to be column vectors unless specifically indicated otherwise. The transpose of a vector or matrix will be denoted by the use of a superscript $'$. Inequalities between vectors in \mathbb{R}^d will be interpreted

componentwise, that is, for $x, y \in \mathbb{R}^d$, $x \leq y$ is equivalent to $x_i \leq y_i$ for $i = 1, \dots, d$. For each $x \in \mathbb{R}^d$ and each set $S \subset \mathbb{R}^d$, the distance between x and S is denoted by

$$d(x, S) = \inf\{|x - y| : y \in S\}.$$

For $x, y \in \mathbb{R}$, $x \vee y = \max\{x, y\}$. For each $x \in \mathbb{R}$, $\lfloor x \rfloor$ denotes the largest integer less than or equal to x . Given a vector $x \in \mathbb{R}^d$, the $d \times d$ diagonal matrix with the entries of x on its diagonal will be denoted by $\text{diag}(x)$. For positive integers d_1 and d_2 , the norm of a $d_1 \times d_2$ matrix A will be given by

$$\|A\| = \left(\sum_{i=1}^{d_1} \sum_{j=1}^{d_2} A_{ij}^2 \right)^{1/2}.$$

The set of non-negative integers will be denoted by \mathbb{Z}_+ and the set of points in \mathbb{R}_+^d with all integer coordinates will be denoted by \mathbb{Z}_+^d . A sum over an empty set of indices will be taken to have a value of zero. The cardinality of a finite set S will be denoted by $|S|$. For $0 \leq s < t < \infty$, any integer $d \geq 1$, and any bounded function $x : [s, t] \rightarrow \mathbb{R}^d$, let $\|x(\cdot)\|_{[s,t]} = \sup_{u \in [s,t]} |x(u)|$ and when $s = 0$, let $\|x(\cdot)\|_t = \|x(\cdot)\|_{[0,t]}$.

All stochastic processes in this paper will be assumed to have sample paths that are right continuous with finite left limits (r.c.l.l.). We denote by $\mathbb{D}([0, \infty), \mathbb{R}^d)$ the space of r.c.l.l. functions from $[0, \infty)$ into \mathbb{R}^d and we endow this space with the usual Skorokhod J_1 -topology. We denote by $\mathbb{C}([0, \infty), \mathbb{R}^d)$ the space of continuous functions from $[0, \infty)$ into \mathbb{R}^d . The Borel σ -algebra on either $\mathbb{D}([0, \infty), \mathbb{R}^d)$ or $\mathbb{C}([0, \infty), \mathbb{R}^d)$ will be denoted by \mathcal{B}^d . Consider X, X^1, X^2, \dots , each of which is a d -dimensional process (possibly defined on different probability spaces). The sequence $\{X^n\}_{n=1}^\infty$ is said to be *tight* if the probability measures induced by the X^n on the measurable space $(\mathbb{D}([0, \infty), \mathbb{R}^d), \mathcal{B}^d)$ form a tight sequence, i.e., they form a weakly relatively compact sequence in the space of probability measures on $(\mathbb{D}([0, \infty), \mathbb{R}^d), \mathcal{B}^d)$. The notation " $X^n \Rightarrow X$ " will mean that, as $n \rightarrow \infty$, the sequence of probability measures induced on $(\mathbb{D}([0, \infty), \mathbb{R}^d), \mathcal{B}^d)$ by $\{X^n\}$ converges weakly to the probability measure induced on the same space by X . We shall describe this in words by saying that X^n converges weakly (or in distribution) to X as $n \rightarrow \infty$. The sequence of processes $\{X^n\}_{n=1}^\infty$ is called *C-tight* if it is tight, and if each weak limit point, obtained as a weak limit along a subsequence, almost surely has sample paths in $\mathbb{C}([0, \infty), \mathbb{R}^d)$.

2. FLOW-LEVEL MODEL.

2.1. Network Structure. We consider a network with finitely many *resources* labeled by $j \in \mathbb{J} \neq \emptyset$. A *route* i is a non-empty subset of \mathbb{J} (interpreted as the set of resources used by route i). We are given a finite, non-empty set \mathbb{I} of allowed

routes. Let $\mathbf{J} = |\mathbb{J}|$, the total number of resources, and $\mathbf{I} = |\mathbb{I}|$, the total number of routes. Let A be the $\mathbf{J} \times \mathbf{I}$ *incidence matrix* which contains only zeros and ones and is defined such that $A_{ji} = 1$ if resource j is used by route i , and $A_{ji} = 0$ otherwise. We assume that A has rank \mathbf{J} , so that it has full row rank. We further assume that resource (bandwidth) *capacities* $(C_j : j \in \mathbb{J})$ are given and that these are all strictly positive and finite.

2.2. Stochastic Primitives. An active flow on route i corresponds to the continuous transmission of a document through the resources used by route i . Transmission is assumed to occur simultaneously through all resources on route i . It is assumed that a new document arrives to route i at each jump time of a Poisson process that has rate parameter $v_i > 0$ and that each such document has an exponentially distributed size with mean $1/\mu_i$ where $\mu_i \in (0, \infty)$. These document sizes are assumed to be independent of one another and to be independent of all arrival times of documents. The number of documents on route i at time zero is assumed to be independent of the remaining sizes of those documents and these sizes are assumed to be independent and exponentially distributed with mean $1/\mu_i$. Initial numbers and sizes of documents, arrival times of new documents and their sizes for different routes $i \in \mathbb{I}$ are assumed to be mutually independent.

2.3. Bandwidth Sharing Policy. Bandwidth capacity is allocated dynamically to the routes according to the following bandwidth sharing policy which was first introduced by Mo and Walrand [35]. The bandwidth for a route is shared equally amongst all of the documents currently being transmitted over that route. Given a fixed parameter $\alpha \in (0, \infty)$ and strictly positive weights $(\kappa_i : i \in \mathbb{I})$, if $N_i(t)$ denotes the (random) number of flows on route i at time t for each $i \in \mathbb{I}$, and $N(t) = (N_i(t) : i \in \mathbb{I})$, then the bandwidth allocated to route i at time t is given by $\Lambda_i(N(t))$ and this bandwidth is shared equally amongst all of the flows on route i , where the function $\Lambda(\cdot) = (\Lambda_i(\cdot) : i \in \mathbb{I})$ is defined as follows (we define it on all of $\mathbb{R}_+^{\mathbf{I}}$ as we shall later apply it to rescaled versions of N).

Let $\Lambda : \mathbb{R}_+^{\mathbf{I}} \rightarrow \mathbb{R}_+^{\mathbf{I}}$ be defined such that for each $n \in \mathbb{R}_+^{\mathbf{I}}$, $\Lambda_i(n) = 0$ for $i \in \mathbb{I}_0(n) \equiv \{l \in \mathbb{I} : n_l = 0\}$, and when $\mathbb{I}_+(n) \equiv \{l \in \mathbb{I} : n_l > 0\}$ is non-empty, $\Lambda^+(n) \equiv (\Lambda_i(n) : i \in \mathbb{I}_+(n))$ is the unique value of $\Lambda^+ = (\Lambda_i : i \in \mathbb{I}_+(n))$ that solves the optimization problem:

$$(1) \quad \begin{array}{ll} \text{maximize} & G_n(\Lambda^+) \\ \text{subject to} & \sum_{i \in \mathbb{I}_+(n)} A_{ji} \Lambda_i \leq C_j, \quad j \in \mathbb{J}, \\ \text{over} & \Lambda_i \geq 0, \quad i \in \mathbb{I}_+(n), \end{array}$$

where for $n \in \mathbb{R}_+^{\mathbf{I}} \setminus \{0\}$ and $\Lambda^+ = (\Lambda_i : i \in \mathbb{I}_+(n)) \in \mathbb{R}_+^{|\mathbb{I}_+(n)|}$,

$$(2) \quad G_n(\Lambda^+) = \begin{cases} \sum_{i \in \mathbb{I}_+(n)} \kappa_i n_i^\alpha \frac{\Lambda_i^{1-\alpha}}{1-\alpha} & \text{if } \alpha \neq 1, \\ \sum_{i \in \mathbb{I}_+(n)} \kappa_i n_i \log \Lambda_i & \text{if } \alpha = 1, \end{cases}$$

and the value of the right member above is taken to be $-\infty$ if $\alpha \in [1, \infty)$ and $\Lambda_i = 0$ for some $i \in \mathbb{I}_+(n)$. The resulting bandwidth allocation is called a *weighted α -fair allocation*.

The properties of the function Λ are summarized in the following proposition. This proposition is proved in the Appendix of Kelly and Williams [26].

PROPOSITION 2.1. *For each $n \in \mathbb{R}_+^{\mathbf{I}}$,*

- (i) $\Lambda_i(n) > 0$ for each $i \in \mathbb{I}_+(n)$,
- (ii) $\Lambda(rn) = \Lambda(n)$ for each $r > 0$,
- (iii) $\Lambda_i(\cdot)$ is continuous at n for those i such that $n_i > 0$, and
- (iv) there exists at least one $p \in \mathbb{R}_+^{\mathbf{J}}$, depending on n , such that

$$(3) \quad \Lambda_i(n) = n_i \left(\frac{\kappa_i}{\sum_{j \in \mathbb{J}} p_j A_{ji}} \right)^{1/\alpha}, \quad \text{for all } i \in \mathbb{I}_+(n),$$

where

$$(4) \quad p_j \left(C_j - \sum_{i \in \mathbb{I}} A_{ji} \Lambda_i(n) \right) = 0 \quad \text{for all } j \in \mathbb{J}.$$

The $(p_j : j \in \mathbb{J})$ are Lagrange multipliers (or *dual variables*) for the optimization problem, where there is one multiplier for each of the capacity constraints.

2.4. *Stochastic Process Description.* The flow count process $N = (N_i : i \in \mathbb{I})$ is a Markov process with state space $\mathbb{Z}_+^{\mathbf{I}}$. We use the following (equivalent in distribution) representation for N and the cumulative unused capacity process $U = (U_j : j \in \mathbb{J})$:

$$(5) \quad N_i(t) = N_i(0) + E_i(t) - S_i(T_i(t)), \quad i \in \mathbb{I},$$

$$(6) \quad U_j(t) = C_j t - \sum_{i \in \mathbb{I}} A_{ji} T_i(t), \quad j \in \mathbb{J},$$

where E_i is a Poisson process with rate ν_i , S_i is a Poisson process with rate μ_i , and $T_i(t)$ is the cumulative amount of bandwidth allocated to route i up to time t and

$$(7) \quad T_i(t) = \int_0^t \Lambda_i(N(s)) ds.$$

We assume that for each $i \in \mathbb{I}$, E_i and S_i are represented by

$$(8) \quad E_i(t) = \sup \left\{ n \geq 0 : \sum_{l=1}^n \xi_i(l) \leq t \right\},$$

and

$$(9) \quad S_i(t) = \sup \left\{ n \geq 0 : \sum_{l=1}^n \zeta_i(l) \leq t \right\},$$

where $\{\xi_i(l)\}_{l=1}^{\infty}$ is a sequence of i.i.d. exponential random variables with mean $1/v_i$ and $\{\zeta_i(l)\}_{l=1}^{\infty}$ is a sequence of i.i.d. exponential random variables with mean $1/\mu_i$. It is assumed that $\{\xi_i(l)\}_{l=1}^{\infty}$, $\{\zeta_i(l)\}_{l=1}^{\infty}$, $N_i(0)$, for $i \in \mathbb{I}$, are mutually independent. We define an (average) *workload process* by

$$(10) \quad W(t) = AM^{-1}N(t) \quad \text{for all } t \geq 0,$$

where $M = \text{diag}(\mu)$ is the $\mathbf{I} \times \mathbf{I}$ diagonal matrix with the entries of μ on its diagonal.

3. SEQUENCE OF SYSTEMS AND SCALING. Consider an increasing sequence of positive scale parameters $\{r_l\}_{l=1}^{\infty}$ which converges to infinity. To ease the notation, we shall simply write r in place of r_l , where it is understood that r increases to infinity through a sequence. We consider a sequence of flow level models indexed by r where the network structure with parameters A and C and bandwidth sharing policy with parameters α and $\{\kappa_i, i \in \mathbb{I}\}$ do not vary with r . Each member of the sequence is a stochastic system as described in the previous section. We append a superscript of r to any process, sequence of random variables or parameter associated with the r^{th} system that depends on r . Thus, we have processes $N^r, W^r, U^r, T^r, E^r, S^r$, sequences of random variables $\xi_i^r = \{\xi_i^r(l)\}_{l=1}^{\infty}$ and $\zeta_i^r = \{\zeta_i^r(l)\}_{l=1}^{\infty}$ for $i \in \mathbb{I}$, parameters v^r , and μ^r , and matrices M^r . Let $\rho_i^r = v_i^r/\mu_i^r$ for each $i \in \mathbb{I}$. We shall assume henceforth that the following *heavy traffic* condition holds.

ASSUMPTION 3.1. (*Heavy Traffic*) *There are $v, \mu \in \mathbb{R}_+^{\mathbf{I}}$ and $\theta \in \mathbb{R}^{\mathbf{J}}$ such that $v_i > 0$ and $\mu_i > 0$ for all $i \in \mathbb{I}$,*

$$(11) \quad v^r \rightarrow v \quad \text{and} \quad \mu^r \rightarrow \mu \quad \text{as } r \rightarrow \infty,$$

$$(12) \quad r(A\rho^r - C) \rightarrow \theta \quad \text{as } r \rightarrow \infty.$$

Let $M = \text{diag}(\mu)$ and $\rho_i = \frac{v_i}{\mu_i}$ for all $i \in \mathbb{I}$. We note that (11)–(12) imply that $\rho^r \rightarrow \rho$ as $r \rightarrow \infty$ and $A\rho = C$.

REMARK 3.1. *Assumption 3.1 thus implies that all resources are heavily loaded. We do not consider the case where some resources are underloaded, however, we conjecture that the diffusion approximation in this case would be as if these underloaded resources were deleted from the model.*

We define fluid scaled processes $\bar{N}^r, \bar{W}^r, \bar{U}^r, \bar{T}^r, \bar{E}^r, \bar{S}^r$ as follows. For each r and $t \geq 0$, let

$$(13) \quad \bar{N}^r(t) = N^r(rt)/r, \quad \bar{W}^r(t) = W^r(rt)/r,$$

$$(14) \quad \bar{U}^r(t) = U^r(rt)/r, \quad \bar{T}^r(t) = T^r(rt)/r,$$

$$(15) \quad \bar{E}^r(t) = E^r(rt)/r, \quad \bar{S}^r(t) = S^r(rt)/r.$$

We define diffusion scaled processes $\hat{N}^r, \hat{W}^r, \hat{U}^r, \hat{E}^r, \hat{S}^r$ as follows. For each r and $t \geq 0$, let

$$(16) \quad \hat{N}^r(t) = \frac{N^r(r^2t)}{r},$$

$$(17) \quad \hat{W}^r(t) = \frac{W^r(r^2t)}{r} = A(M^r)^{-1}\hat{N}^r(t),$$

$$(18) \quad \hat{U}^r(t) = \frac{U^r(r^2t)}{r},$$

$$(19) \quad \hat{E}^r(t) = \frac{E^r(r^2t) - \nu^r r^2t}{r},$$

$$(20) \quad \hat{S}^r(t) = \frac{S^r(r^2t) - \mu^r r^2t}{r}.$$

As \hat{E}_i^r, \hat{S}_i^r , $i \in \mathbb{I}$, are independent Poisson processes with parameters satisfying the convergence conditions (11), it follows that we have the following well known functional central limit result [2]:

$$(21) \quad (\hat{E}^r, \hat{S}^r) \Rightarrow (\tilde{E}, \tilde{S}), \quad \text{as } r \rightarrow \infty,$$

where \tilde{E} and \tilde{S} are independent \mathbf{I} -dimensional Brownian motions starting from the origin with zero drift and covariance matrices $\text{diag}(\nu)$ and $\text{diag}(\mu)$, respectively.

Finally, we assume that independent of (21), $\hat{W}^r(0)$ converges in distribution as $r \rightarrow \infty$ to a \mathbf{J} -dimensional random variable.

4. FLUID MODEL. In this section we recall some definitions and results established in the prior work [26]. These will be needed for our statement and proof of multiplicative state space collapse.

4.1. *Fluid Model Solution.* A fluid model solution can be thought of as a formal limit of the sequence $\{\bar{N}^r\}$ as $r \rightarrow \infty$. In fact, if one assumes that $\bar{N}^r(0)$ converges in distribution as $r \rightarrow \infty$ to a random variable taking values in $\mathbb{R}_+^{\mathbf{I}}$, then one can show (cf. Appendix to [26]) that $\{(\bar{T}^r, \bar{E}^r, \bar{S}^r, \bar{N}^r, \bar{U}^r)\}$ is C -tight and any weak limit point $(\bar{T}, \bar{E}, \bar{S}, \bar{N}, \bar{U})$ yields a fluid model solution \bar{N} a.s.

The following notions are used in the definition of a fluid model solution given below. A function $f = (f_1, \dots, f_{\mathbf{I}}) : [0, \infty) \rightarrow \mathbb{R}_+^{\mathbf{I}}$ is absolutely continuous if each of its components $f_i : [0, \infty) \rightarrow \mathbb{R}_+$, $i = 1, \dots, \mathbf{I}$, is absolutely continuous. A *regular point* for an absolutely continuous function $f : [0, \infty) \rightarrow \mathbb{R}_+^{\mathbf{I}}$ is a value of $t \in (0, \infty)$ at which each component of f is differentiable. (Since f is absolutely continuous, almost every time $t \in (0, \infty)$ is a regular point for f . Furthermore, f can be recovered by integration from its a.e. defined derivative.)

DEFINITION 4.1. *A fluid model solution is an absolutely continuous function $n : [0, \infty) \rightarrow \mathbb{R}_+^{\mathbf{I}}$ such that at each regular point $t > 0$ for $n(\cdot)$ we have for each $i \in \mathbb{I}$,*

$$(22) \quad \frac{d}{dt} n_i(t) = \begin{cases} v_i - \mu_i \Lambda_i(n(t)), & \text{if } n_i(t) > 0, \\ 0, & \text{if } n_i(t) = 0, \end{cases}$$

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

$$(23) \quad \sum_{i \in \mathbb{I}_+(n(t))} A_{ji} \Lambda_i(n(t)) + \sum_{i \in \mathbb{I}_0(n(t))} A_{ji} \rho_i \leq C_j,$$

where $\mathbb{I}_+(n(t)) = \{i \in \mathbb{I} : n_i(t) > 0\}$ and $\mathbb{I}_0(n(t)) = \{i \in \mathbb{I} : n_i(t) = 0\}$.

REMARK 4.1. *Note that we are not assuming uniqueness of fluid model solutions given the initial state.*

4.2. Invariant Manifold.

DEFINITION 4.2. *A state $n_0 \in \mathbb{R}_+^{\mathbf{I}}$ is called invariant (for the fluid model) if there is a fluid model solution $n(\cdot)$ such that $n(t) = n_0$ for all $t \geq 0$. Let \mathcal{M}_α denote the set of all invariant states. We call \mathcal{M}_α the invariant manifold.*

REMARK 4.2. *Although $\alpha \in (0, \infty)$ is fixed throughout, we indicate the dependence of \mathcal{M}_α on α explicitly here as it will be useful later on when we explain how the state space for the proposed workload diffusion approximation varies with α .*

Various characterizations of the invariant states were given in [26]. We summarize these in Theorem 4.1 below. For this, we need the following definitions.

For each $n \in \mathbb{R}_+^{\mathbf{I}}$, define $w(n) = (w_j(n) : j \in \mathbb{J})$ to be given by

$$(24) \quad w_j(n) = \sum_{i \in \mathbb{I}} A_{ji} \frac{n_i}{\mu_i}, \quad j \in \mathbb{J}.$$

We call $w(n)$ the *workload* associated with n .

For each $w \in \mathbb{R}_+^{\mathbf{J}}$, define $\Delta(w)$ to be the unique value of $n \in \mathbb{R}_+^{\mathbf{I}}$ that solves the following optimization problem:

$$(25) \quad \begin{array}{ll} \text{minimize} & F(n) \\ \text{subject to} & \sum_{i \in \mathbb{I}} A_{ji} \frac{n_i}{\mu_i} \geq w_j, \quad j \in \mathbb{J}, \\ \text{over} & n_i \geq 0, \quad i \in \mathbb{I}, \end{array}$$

where

$$(26) \quad F(n) = \frac{1}{\alpha + 1} \sum_{i \in \mathbb{I}} v_i \kappa_i \mu_i^{\alpha-1} \left(\frac{n_i}{v_i} \right)^{\alpha+1}, \quad n \in \mathbb{R}_+^{\mathbf{I}}.$$

(This function F was introduced in [3] as a Lyapunov function for the fluid model. In fact, it is a Lyapunov function for the original flow count process N and can be used to show positive recurrence of N when the average load on each resource is less than its capacity.) The function Δ has the two properties stated in the next proposition.

PROPOSITION 4.1. *The function $\Delta : \mathbb{R}_+^{\mathbf{J}} \rightarrow \mathbb{R}_+^{\mathbf{I}}$ is continuous. Furthermore, for each $w \in \mathbb{R}_+^{\mathbf{J}}$ and $c > 0$,*

$$(27) \quad \Delta(cw) = c\Delta(w).$$

PROOF. The first property is proved in Lemma 6.3 of [26]. For the second property, note that for $w \in \mathbb{R}_+^{\mathbf{J}}$ and $c > 0$, $c\Delta(w)$ satisfies the constraints in (25) with cw in place of w and so

$$(28) \quad c^{\alpha+1} F(\Delta(w)) = F(c\Delta(w)) \geq F(\Delta(cw)).$$

On the other hand, by writing w/c in place of w in (28), we find that

$$c^{\alpha+1} F(\Delta(w/c)) \geq F(\Delta(w)),$$

and then by replacing c by $1/c$ and rearranging, we obtain

$$(29) \quad F(\Delta(cw)) \geq c^{\alpha+1} F(\Delta(w)).$$

On combining (28) and (29), we conclude that

$$F(\Delta(cw)) = F(c\Delta(w))$$

and by uniqueness of the solution to (25), we obtain the second property. \square

THEOREM 4.1. *The following are equivalent for $n \in \mathbb{R}_+^{\mathbb{I}}$.*

- (i) n is an invariant state (for the fluid model), i.e., $n \in \mathcal{M}_\alpha$,
- (ii) $\Lambda_i(n) = \rho_i$ for all $i \in \mathbb{I}_+(n) = \{l \in \mathbb{I} : n_l > 0\}$,
- (iii) there is $q \in \mathbb{R}_+^{\mathbb{J}}$ such that

$$(30) \quad n_i = \rho_i \left(\frac{\sum_{j \in \mathbb{J}} q_j A_{ji}}{\kappa_i} \right)^{1/\alpha} \quad \text{for all } i \in \mathbb{I},$$

- (iv) $n = \Delta(w(n))$.

PROOF. This follows immediately from Lemma 5.1 and Theorems 5.1, 5.3 of [26]. \square

REMARK 4.3. *Note that if the conditions of Theorem 4.1 are satisfied, then $p = q$ satisfies conditions (3) and (4) and thus $(q_j : j \in \mathbb{J})$ are dual variables for the optimization problem (1); note that we use the fact that $A\rho = C$ for this. We have chosen to use q to denote the dual variables associated with the invariant states to distinguish them from the dual variables p associated with arbitrary states n . This distinction will be useful in our proof of convergence to a diffusion process (cf. Lemma 7.5), where we need to distinguish the dual variables associated with actual system states from the dual variables associated with nearby points on the invariant manifold. It is important to make this distinction because when a system state is near an invariant state and some component of the system state is near zero, it need not follow that the dual variables associated with the two states are close.*

PROPOSITION 4.2. *For each $w \in \mathbb{R}_+^{\mathbb{J}}$, $\Delta(w) \in \mathcal{M}_\alpha$, i.e., $\Delta(w)$ is an invariant state.*

PROOF. Let $w \in \mathbb{R}_+^{\mathbb{J}}$. Since $\Delta(w)$ is the unique optimal solution to (25), it follows from Lemma 6.4 of Kelly and Williams [26] that there is $q \in \mathbb{R}_+^{\mathbb{J}}$ such that $\Delta(w)_i = \rho_i \left(\frac{\sum_{j \in \mathbb{J}} q_j A_{ji}}{\kappa_i} \right)^{1/\alpha}$ for all $i \in \mathbb{I}$. Then by Theorem 4.1, we have that $\Delta(w)$ is an invariant state. \square

4.3. *Asymptotic Properties of Fluid Model Solutions.* The next three propositions note some properties of fluid model solutions that follow from the analysis in [26] and that are used in our proof of multiplicative state space collapse (cf. Theorem 5.1 below).

PROPOSITION 4.3. *For each $R \in (0, \infty)$, there is a constant $D(R) \in [R, \infty)$ such that for any fluid model solution $n(\cdot)$ satisfying $|n(0)| \leq R$, we have $|n(t)| \leq D(R)$ for all $t \geq 0$.*

PROOF. The proof of this proposition is implicit in the proof of Theorem 5.2 in [26]. \square

The next proposition states that fluid model solutions converge uniformly to the invariant manifold \mathcal{M}_α , where the uniformity applies across all fluid model solutions that start inside a compact subset of $\mathbb{R}_+^{\mathbf{I}}$.

PROPOSITION 4.4. *Fix $R \in (0, \infty)$ and $\varepsilon > 0$. There is a constant $T_{R,\varepsilon} \in [1, \infty)$ such that for each fluid model solution $n(\cdot)$ satisfying $|n(0)| \leq R$ we have*

$$(31) \quad d(n(t), \mathcal{M}_\alpha) < \varepsilon \quad \text{for all } t \geq T_{R,\varepsilon}.$$

PROOF. The content of this proposition is the same as that of Theorem 5.2 in [26]. \square

PROPOSITION 4.5. *For each $R \in (0, \infty)$ and $\varepsilon > 0$, there is $\delta > 0$ such that for any fluid model solution $n(\cdot)$ satisfying $|n(0)| \leq R$ and $d(n(0), \mathcal{M}_\alpha) < \delta$, we have $d(n(t), \mathcal{M}_\alpha) < \varepsilon$ for all $t \geq 0$.*

PROOF. This proposition follows from the proof of Theorem 5.2 in [26]. For completeness, we provide a few details. Fix $R > 0$ and $\varepsilon > 0$. By Proposition 4.3, there is a compact set $B(R)$ in $\mathbb{R}_+^{\mathbf{I}}$ such that $n(t) \in B(R)$ for all $t \geq 0$ for any fluid model solution n satisfying $|n(0)| \leq R$. Let $D = \{u \in B(R) : d(u, \mathcal{M}_\alpha) \geq \varepsilon\}$. As shown in [26], there is a continuous function $H : \mathbb{R}_+^{\mathbf{I}} \rightarrow \mathbb{R}_+$ that is zero on \mathcal{M}_α and strictly positive off \mathcal{M}_α such that $H(n(\cdot))$ is non-increasing for each fluid model solution n . Let $\delta_1 = \inf\{H(u) : u \in D\}$. Then, $\delta_1 > 0$ and by the properties of H , there is $\delta > 0$ such that whenever n is a fluid model solution satisfying $|n(0)| \leq R$ and $d(n(0), \mathcal{M}_\alpha) \leq \delta$, then $H(n(0)) < \delta_1$, and since $H(n(\cdot))$ is a non-increasing function, it follows that $H(n(t)) < \delta_1$ for all $t \geq 0$. The latter implies that $n(t) \notin D$ for all $t \geq 0$ and so $d(n(t), \mathcal{M}_\alpha) < \varepsilon$ for all $t \geq 0$. \square

The following corollary shows that fluid model solutions starting on the invariant manifold stay at their starting points for all time. We note this for the reader's interest. We do not use this corollary in our proofs.

COROLLARY 4.1. *Suppose that $n(\cdot)$ is a fluid model solution such that $n(0) \in \mathcal{M}_\alpha$. Then $n(t) = n(0) \in \mathcal{M}_\alpha$ for all $t \geq 0$.*

PROOF. By Proposition 4.5, since $d(n(0), \mathcal{M}_\alpha) = 0$, then $d(n(t), \mathcal{M}_\alpha) = 0$ for all $t \geq 0$ and hence $n(t) \in \mathcal{M}_\alpha$ for all $t \geq 0$. It follows from Theorem 4.1 that for

each $t > 0$, $\Lambda_i(n(t)) = \rho_i = v_i/\mu_i$ for all i such that $n_i(t) > 0$. Then by the fluid model dynamics (22), for each regular point $t > 0$ of $n(\cdot)$, we have

$$(32) \quad \frac{d}{dt}n_i(t) = v_i - \mu_i\rho_i = 0,$$

if $n_i(t) > 0$, and the last equality also holds if $n_i(t) = 0$. It follows, since the absolutely continuous function $n(\cdot)$ can be recovered from its almost everywhere defined derivative, that $n(t) = n(0)$ for all $t \geq 0$. \square

5. MAIN RESULTS. In this section, we describe the main results of this paper. We begin with our result on *multiplicative state space collapse*. This is established using the asymptotic behavior of fluid model solutions described in Section 4.3. Loosely speaking, multiplicative state space collapse shows that an approximation for \hat{N}^r can be derived from one for \hat{W}^r via the continuous lifting map Δ (see (25)–(26) for the definition of this map). This lifting map can be nonlinear (for $\alpha \neq 1$). The multiplicative state space collapse result leads to a natural conjecture for a diffusion approximation to \hat{W}^r . In the case $\alpha = 1$, assuming a mild local traffic condition and suitable initial conditions, we prove that the conjectured diffusion approximation is valid. (In Section 5.6, we indicate some of the challenges associated with establishing this conjecture for $\alpha \neq 1$.) When $\alpha = 1$ and all of the weights for the bandwidth sharing policy are equal (proportional fair sharing), we use results of Harrison and Williams [17] and Williams [41] to show that the diffusion has a product form invariant measure. When this measure has finite total mass, this result suggests an approximation for the stationary distribution of the flow count process which we are able to extend to the case where the document size distributions are finite mixtures of exponential distributions and to some models with multi-path routing. So as not to disrupt the flow of results and associated discussion, we defer the longish proofs of multiplicative state space collapse and of the diffusion approximation to Sections 6 and 7, respectively.

5.1. Multiplicative State Space Collapse.

DEFINITION 5.1. (*Multiplicative State Space Collapse*) *Multiplicative state space collapse holds (for the sequence of flow level models described in Section 3), if for each $T > 0$,*

$$(33) \quad \frac{\|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T}{\|\hat{N}^r(\cdot)\|_T \vee 1} \rightarrow 0$$

in probability as $r \rightarrow \infty$.

REMARK 5.1. *We note here that in our form of multiplicative state space collapse, the normalization (in the denominator) is in terms of the flow count process, whereas in Bramson's version for multiclass queueing networks [5], it is in terms of a workload process. Furthermore, the lifting maps in [5] are all linear whereas here Δ can be nonlinear (for $\alpha \neq 1$).*

REMARK 5.2. *If (33) holds without the factor in the denominator, then state space collapse is said to hold. Multiplicative state space collapse is more convenient for the purpose of verification and if $\{\hat{N}^r\}$ (or $\{\hat{W}^r\}$) satisfies a compact containment condition then state space collapse follows from multiplicative state space collapse. As was the case for open multiclass HL queueing networks [42], in establishing our diffusion approximation result for $\alpha = 1$ under a mild local traffic condition, we will show for this case that multiplicative state space collapse implies state space collapse.*

The following theorem is one of the main results of this paper. It is proved in Section 6.

THEOREM 5.1. *Assume that*

$$(34) \quad |\hat{N}^r(0) - \Delta(\hat{W}^r(0))| \rightarrow 0$$

in probability as $r \rightarrow \infty$. Then multiplicative state space collapse holds.

5.2. *Conjectured Diffusion Approximation.* We are interested in obtaining a diffusion approximation for the scaled workload process \hat{W}^r . The multiplicative state space collapse result can then be used to obtain a diffusion approximation for the scaled flow count process \hat{N}^r .

For each r , define the double fluid scaled bandwidth allocation process

$$(35) \quad \bar{T}^r(t) = \frac{T^r(r^2 t)}{r^2}, \quad t \geq 0.$$

Using (5)–(7) and (10) for the r^{th} system, the definitions of rescaled processes, and (ii) of Proposition 2.1, after some simple manipulations, we obtain for all $t \geq 0$:

$$(36) \quad \hat{W}^r(t) = \hat{W}^r(0) + \hat{X}^r(t) + \hat{U}^r(t),$$

where

$$(37) \quad \hat{X}^r(t) = A(M^r)^{-1} \left(\hat{E}^r(t) - \hat{S}^r(\bar{T}^r(t)) \right) + r(A\rho^r - C)t,$$

$$(38) \quad (\hat{S}^r(\bar{T}^r(t)))_i = \hat{S}_i^r(\bar{T}_i^r(t)), \quad i \in \mathbb{I},$$

$$(39) \quad \begin{aligned} \hat{U}^r(t) &= r(Ct - A\bar{T}^r(t)) \\ &= r^{-1} \int_0^{r^2 t} (C - A\Lambda(N^r(s))) ds \\ &= r \int_0^t (C - A\Lambda(\hat{N}^r(s))) ds. \end{aligned}$$

If we postulate that multiplicative state space collapse implies state space collapse, then by formally (non-rigorously) passing to the limit in the expression (36) for \hat{W}^r , we can obtain a natural conjecture for a diffusion approximation to \hat{W}^r . Immediately below, we give an informal description of how one might arrive at this conjecture. Following that we give a precise mathematical description of the diffusion process and of the conjecture.

For the following informal description, which is used to motivate the form of the conjectured diffusion approximation, we postulate that the sequence of processes $\{(\hat{W}^r, \bar{T}^r, \hat{U}^r, \hat{E}^r, \hat{S}^r)\}$ converges in distribution to a 5-tuple of continuous processes $(\tilde{W}, T^*, \tilde{U}, \tilde{E}, \tilde{S})$. We also postulate that state space collapse (SSC) holds (not just multiplicative state space collapse). From the convergence of $\{\hat{W}^r\}$, SSC and continuity of the lifting map Δ , it follows that \hat{N}^r converges in distribution to a continuous process $\tilde{N} = \Delta(\tilde{W})$ that lives on the invariant manifold \mathcal{M}_α . The fact that $\hat{W}^r = A(M^r)^{-1}\hat{N}^r$ will yield in the limit that $\tilde{W} = AM^{-1}\tilde{N}$. By the characterization of invariant states given in Theorem 4.1, it will then follow that for each $t \geq 0$ and realization ω , there is $q(t, \omega) \in \mathbb{R}_+^{\mathbf{J}}$ such that

$$(40) \quad \tilde{N}_i(t, \omega) = \rho_i \left(\frac{\sum_{j \in \mathbb{J}} q_j(t, \omega) A_{ji}}{\kappa_i} \right)^{1/\alpha}, \quad \text{for all } i \in \mathbb{I}.$$

Consequently, $\tilde{W} = AM^{-1}\tilde{N}$ will live in the space

$$(41) \quad \mathcal{W}_\alpha = AM^{-1}\mathcal{M}_\alpha$$

where

$$(42) \quad \mathcal{M}_\alpha = \left\{ n \in \mathbb{R}_+^{\mathbf{I}} : n_i = \rho_i \left(\frac{(q'A)_i}{\kappa_i} \right)^{\frac{1}{\alpha}} \text{ for all } i \in \mathbb{I}, \text{ some } q \in \mathbb{R}_+^{\mathbf{J}} \right\}.$$

We call \mathcal{W}_α the *workload cone*. This is the state space for the conjectured diffusion approximation \tilde{W} .

By the assumption made at the end of Section 3, we know that $\hat{W}^r(0)$ converges in distribution, independently of the primitive arrival and service processes. We denote the limit distribution of $\hat{W}^r(0)$ by η . Under the postulated convergence of \hat{W}^r , η will be concentrated on \mathscr{W}_α .

We now turn our attention to the term \hat{X}^r in the expression (36) for \hat{W}^r . Given the functional central limit theorem result (21) for the diffusion scaled arrival and service processes (\hat{E}^r, \hat{S}^r) , and the heavy traffic Assumption 3.1, if we postulate that the double fluid scaled allocation processes \tilde{T}^r achieve the nominal levels given by $T^*(t) \equiv \rho t$ in the heavy traffic limit, then \hat{X}^r given by (37) will converge in distribution to the Brownian motion $AM^{-1}(\tilde{E}(\cdot) - \tilde{S}(T^*(\cdot))) + \theta(\cdot)$, where $(\tilde{S}(T^*(\cdot)))_i = \tilde{S}_i(T_i^*(\cdot))$ for $i \in \mathbb{I}$, $\theta(t) = \theta t$ for all $t \geq 0$ and θ is defined in the heavy traffic Assumption 3.1. This Brownian motion starts from the origin, it has drift θ and covariance matrix $AM^{-1}\text{diag}(v + v)M^{-1}A'$, where we have used the facts that $M = \text{diag}(\mu)$ and $\mu_i \rho_i = v_i$ for all $i \in \mathbb{I}$ to compute the second term in the diagonal part of the covariance matrix expression.

On examining the representation (36) for \hat{W}^r , we see that it remains to conjecture properties for the postulated limit \tilde{U} of the scaled unused capacity process \hat{U}^r as $r \rightarrow \infty$. The limit \tilde{U} will inherit the non-decreasing property from the \hat{U}^r . The main issue is to determine where each of the components of \tilde{U} can increase. For the prelimit process, \hat{U}^r , to determine where its components increase, we see from (39) that it suffices to identify where each of the components of $C - A\Lambda(\hat{N}^r(\cdot))$ is strictly positive. From Proposition 2.1, if $j \in \mathbb{J}$ such that $C_j - \sum_{i \in \mathbb{I}} A_{ji} \Lambda_i(\hat{N}^r(t, \omega)) > 0$, then there is a Lagrange multiplier $p^r(t, \omega) \in \mathbb{R}_+^{\mathbb{J}}$ such that $p_j^r(t, \omega) = 0$ and

$$(43) \quad \hat{N}_i^r(t, \omega) = \Lambda_i(\hat{N}^r(t, \omega)) \left(\frac{\sum_{k \in \mathbb{J}} p_k^r(t, \omega) A_{ki}}{\kappa_i} \right)^{1/\alpha} \quad \text{for all } i \in \mathbb{I}.$$

(For this, we note that both sides of (43) are zero if $\hat{N}_i^r(t, \omega) = 0$.) If $\Lambda(\hat{N}^r(t, \omega))$ is close to the nominal allocation ρ , then by (43), $\hat{N}^r(t, \omega)$ is near

$$(44) \quad \mathcal{M}_\alpha^j = \left\{ n \in \mathbb{R}_+^{\mathbb{I}} : n_i = \rho_i \left(\frac{(q'A)_i}{\kappa_i} \right)^{\frac{1}{\alpha}} \text{ for all } i \in \mathbb{I}, \right. \\ \left. \text{some } q \in \mathbb{R}_+^{\mathbb{J}} \text{ satisfying } q_j = 0 \right\},$$

the subset of \mathcal{M}_α obtained by setting q_j equal to zero in (42), and then $\hat{W}^r = A(M^r)^{-1}\hat{N}^r$ will be close to

$$(45) \quad \mathscr{W}_\alpha^j = \{AM^{-1}n : n \in \mathcal{M}_\alpha^j\},$$

which we refer to as the j^{th} face of the workload cone \mathscr{W}_α . Thus one might conjecture that in the limit, \tilde{U}_j can increase only when \tilde{W} is on the face \mathscr{W}_α^j .

Putting all of the above musings together leads to the informal conjecture that \hat{W}^r converges in distribution to a \mathbf{J} -dimensional diffusion process \tilde{W} of the form $\tilde{W}(0) + \tilde{X} + \tilde{U}$. The state space for \tilde{W} is the workload cone \mathcal{W}_α and the initial distribution of \tilde{W} is given by η . In the interior of \mathcal{W}_α , the increments of \tilde{W} are given by the increments of a Brownian motion \tilde{X} with drift θ , covariance matrix

$$(46) \quad \Gamma = 2AM^{-1}\text{diag}(v)M^{-1}A',$$

that starts from the origin. The process \tilde{W} is confined to the cone \mathcal{W}_α by instantaneous “pushing” at the boundary of \mathcal{W}_α . The direction of push allowed when \tilde{W} is on the boundary face \mathcal{W}_α^j is γ^j , the unit vector parallel to the positive j^{th} coordinate axis in $\mathbb{R}_+^{\mathbf{J}}$, and the cumulative amount of push in that direction is given by the j^{th} component of the continuous non-decreasing process \tilde{U} . (At the intersections of boundary faces, the combined effect of pushing on the individual boundary faces is a push in the direction of a convex combination of the pushing directions available from the intersecting boundary faces.) The direction of push on a given boundary face is usually called a *direction of reflection*, whereas it might be better thought of as a direction of regulation for the process. (The term reflection comes from the fact that in one-dimension when the drift is zero, the construction of such a regulated process from a Brownian motion can be done by a mirror reflection; although this type of construction does not generally apply in higher dimensions, the term reflection is still used.)

We now introduce a precise definition for the conjectured diffusion approximation to the workload process \hat{W}^r . (The delicate issue of existence and uniqueness for this process is discussed further below.) This definition will be used in giving a precise statement of our conjecture. Here θ is a vector in $\mathbb{R}^{\mathbf{J}}$, Γ is given by (46), γ^j is the unit vector parallel to the positive j^{th} coordinate axis and η is a Borel probability measure on \mathcal{W}_α .

DEFINITION 5.2. *A Semimartingale Reflecting Brownian Motion that lives in the cone \mathcal{W}_α , has direction of reflection γ^j on the boundary face \mathcal{W}_α^j for each $j \in \mathbb{J}$, has drift θ and covariance matrix Γ , and has initial distribution η on \mathcal{W}_α , is an adapted, \mathbf{J} -dimensional process \tilde{W} defined on some filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ such that*

- (i) *P*-a.s., $\tilde{W}(t) = \tilde{W}(0) + \tilde{X}(t) + \tilde{U}(t)$ for all $t \geq 0$,
- (ii) *P*-a.s., \tilde{W} has continuous paths, $\tilde{W}(t) \in \mathcal{W}_\alpha$ for all $t \geq 0$, and $\tilde{W}(0)$ has distribution η ,
- (iii) under *P*,
 - (a) \tilde{X} is a \mathbf{J} -dimensional Brownian motion starting from the origin with drift θ and covariance matrix Γ ,

- (b) $\{\tilde{X}(t) - \theta t, \mathcal{F}_t, t \geq 0\}$ is a martingale,
- (iv) for each $j \in \mathbb{J}$, \tilde{U}_j is an adapted, one-dimensional process such that P-a.s.,
 - (a) $\tilde{U}_j(0) = 0$,
 - (b) \tilde{U}_j is continuous and non-decreasing,
 - (c) $\tilde{U}_j(t) = \int_0^t 1_{\{\tilde{W}(s) \in \mathcal{W}_\alpha^j\}} d\tilde{U}_j(s)$ for all $t \geq 0$.

REMARK 5.3. We call a process \tilde{W} satisfying the above properties an SRBM associated with the data $(\mathcal{W}_\alpha, \theta, \Gamma, \{\gamma^j : j \in \mathbb{J}\}, \eta)$. Here adapted means adapted to the filtration $\{\mathcal{F}_t\}$. We note that this filtration need not be the one generated by \tilde{X} , it can be larger. However, it can always be taken to be the filtration generated by $\tilde{W}, \tilde{X}, \tilde{U}$. The martingale condition on \tilde{X} is included here as we are using a “weak” definition of the process. This martingale property is needed in establishing uniqueness in law for an SRBM. The term “semimartingale” refers to the fact that \tilde{W} is the sum of a continuous martingale and a continuous process that is locally of bounded variation. Condition (iv)(c) corresponds to the condition that \tilde{U}_j can only increase when \tilde{W} is on the boundary face \mathcal{W}_α^j .

We now give a precise statement of our conjecture.

CONJECTURE 5.1. Suppose that the limit distribution of $\hat{W}^r(0)$ is η , a probability distribution on \mathcal{W}_α endowed with the Borel σ -algebra, and suppose that

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

Then, \hat{W}^r converges in distribution as $r \rightarrow \infty$ to a process \tilde{W} that is an SRBM associated with the data $(\mathcal{W}_\alpha, \theta, \Gamma, \{\gamma^j : j \in \mathbb{J}\}, \eta)$.

We shall prove that Conjecture 5.1 holds when $\alpha = 1$, provided a mild local traffic condition holds. We state this result in the next subsection. In the remainder of the current subsection, we indicate some of the challenges associated with giving a rigorous proof of the conjecture.

When $\alpha = 1$ (corresponding to weighted proportional fair sharing), we can express the cone \mathcal{W}_α in the following simple form:

$$(47) \quad \mathcal{W}_1 = \{ABA'q : q \in \mathbb{R}_+^{\mathbf{J}}\},$$

where B is an $\mathbf{I} \times \mathbf{I}$ diagonal matrix with the i^{th} diagonal entry being $\frac{\nu_i}{\mu_i^2 \kappa_i} > 0$. Thus, \mathcal{W}_1 is a polyhedral cone. Since A has full row rank and B is a diagonal matrix with strictly positive diagonal entries, ABA' is a linear bijection between $\mathbb{R}_+^{\mathbf{J}}$ and \mathcal{W}_1 . It follows from this that \mathcal{W}_1 is a simple polyhedral cone. Necessary and

sufficient conditions for the existence and uniqueness in law of SRBMs living in simple polyhedral domains have been given by Dai and Williams [10]. Under these conditions, the SRBM is a diffusion, i.e., a continuous strong Markov process. It will turn out that the conditions of [10] are satisfied by our data when $\alpha = 1$. For $\alpha \neq 1$ and $\mathbf{J} = 2$, the workload cone is a wedge which is still a simple polyhedral cone. However, in general, for $\alpha \neq 1$ and $\mathbf{J} > 2$, the workload cone \mathscr{W}_α is not a polyhedral cone (it has curved boundaries). In this case, we have some partial (unpublished) results on existence and uniqueness for SRBMs. The main impediment to obtaining a general result is that boundary faces can meet in cusp-like singularities making it challenging to even determine whether the process can escape from the cusp and whether it can do so in a unique manner (see Section 5.6 for an example).

Even if one has existence and uniqueness of the SRBM, for any proof of the conjecture, there are a number of other challenges to overcome. Firstly, one needs to establish C -tightness of the sequence of triples $\{(\hat{W}^r, \hat{X}^r, \hat{U}^r)\}$. This is largely an issue of the C -tightness of $\{\hat{U}^r\}$. One also needs to show that multiplicative state space collapse implies state space collapse. One of the most challenging aspects is to show that for any possible limit $(\tilde{W}, \tilde{X}, \tilde{U})$ of the sequence $\{(\hat{W}^r, \hat{X}^r, \hat{U}^r)\}$, for each $j \in \mathbb{J}$, the process \tilde{U}_j can only increase when \tilde{W} is on the boundary face \mathscr{W}_α^j . Indeed, in our informal use of (43) to arrive at our conjecture, we neglected the fact that $\Lambda_i(\hat{N}^r(t, \omega))$ need not be near ρ_i when some component of $\hat{N}^r(t, \omega)$ is near zero (recall that $\Lambda_i(n)$ need not be continuous when n_i is zero), and also the notions of “nearness” and “closeness” used loosely in our informal description are not necessarily uniform. However, if $\hat{N}_i^r(t, \omega)$ is at or near zero when $p_j^r(t, \omega) = 0$ for an i such that $A_{ji} > 0$, then we can show that $\hat{W}^r(t, \omega)$ is near the boundary face \mathscr{W}_α^j . To take advantage of this observation, in the case when we prove the conjecture ($\alpha = 1$), we will assume that a mild local traffic condition holds.

In summary, the main reasons that we are able to treat the case $\alpha = 1$ are that the existence and uniqueness theory for the limit diffusion process is in place [10], and there is an associated invariance principle [20] which loosely speaking is a perturbation result telling us that processes such as \hat{W}^r , that satisfy perturbed versions of the defining conditions for an SRBM, are close in distribution to an SRBM. In particular, for $\alpha = 1$, the invariance principle of [20] takes care of establishing the C -tightness of $\{(\hat{W}^r, \hat{X}^r, \hat{U}^r)\}$ and, in the presence of the uniqueness in law of the SRBM [10], it implies convergence in distribution of \hat{W}^r to an SRBM. In the case $\alpha \neq 1$ and $\mathbf{J} = 2$, \mathscr{W}_α is a wedge (a polyhedral cone) and our proof for $\alpha = 1$ can be extended to this case. In as yet unpublished work, we have been able to establish uniqueness in law of the SRBM and to establish an invariance principle for some cases where $\alpha \neq 1$ and $\mathbf{J} > 2$. However, some cases, especially when boundary faces meet in cusp-like singularities, are as yet unresolved. We summarize the situ-

ation for $\alpha \neq 1$ in Section 5.6. However, because of the partial nature of our results so far, we leave description of these further developments to future work.

5.3. *Diffusion Approximation for Weighted Proportional Fair Sharing ($\alpha = 1$).* The following condition is used in the next theorem. This condition can be interpreted as a *local traffic* assumption under which each resource has at least one route that only uses that resource.

ASSUMPTION 5.1. (*Local Traffic*) For each $j \in \mathbb{J}$, there is at least one $i \in \mathbb{I}$ such that $A_{ji} > 0$ and $A_{ki} = 0$ for all $k \neq j$.

The following theorem is proved in Section 7.

THEOREM 5.2. Assume that $\alpha = 1$ and the local traffic Assumption 5.1 holds. Suppose that the limit distribution of $\hat{W}^r(0)$ as $r \rightarrow \infty$ is η (a probability measure on \mathscr{W}_1), and that $|\hat{N}^r(0) - \Delta(\hat{W}^r(0))| \rightarrow 0$ in probability as $r \rightarrow \infty$. Then (\hat{W}^r, \hat{N}^r) converges in distribution as $r \rightarrow \infty$ to a continuous process (\tilde{W}, \tilde{N}) , where \tilde{W} is an SRBM with data $(\mathscr{W}_1, \theta, \Gamma, \{\gamma^j : j \in \mathbb{J}\}, \eta)$ and $\tilde{N} = \Delta(\tilde{W})$.

In the case $\alpha = 1$, the lifting map Δ is in fact a linear map on \mathscr{W}_1 , given by

$$(48) \quad \Delta(w) = \text{diag}(\rho) \text{diag}(\kappa)^{-1} A' (ABA')^{-1} w, \quad w \in \mathscr{W}_1.$$

Indeed, for $\alpha = 1$ and $w \in \mathscr{W}_1$, if $q = (ABA')^{-1} w$ and n is given by the right hand side of (48), then $w(n) = AM^{-1}n = w$ by the definition of $B = M^{-1} \text{diag}(v) \text{diag}(\kappa)^{-1} M^{-1}$, and $n = \text{diag}(\rho) \text{diag}(\kappa)^{-1} A' q$ so that (30) holds, and then by Theorem 4.1 we conclude that

$$n = \Delta(w(n)) = \Delta(w),$$

and hence (48) holds. By the remark following Theorem 4.1, the $(q_j : j \in \mathbb{J})$ defined above are dual variables for the optimization problem (1). It follows that we can associate a *process of dual random variables* \tilde{Q} with the SRBM \tilde{W} of Theorem 5.2 as follows. Given \tilde{W} as in Theorem 5.2, define

$$(49) \quad \tilde{Q} = (ABA')^{-1} \tilde{W}.$$

This process \tilde{Q} inherits an SRBM structure from \tilde{W} . In fact, \tilde{Q} is a semimartingale reflecting Brownian motion living in $\mathbb{R}_+^{\mathbb{J}}$ having the form:

$$(50) \quad \tilde{Q}(t) = \tilde{Q}(0) + (ABA')^{-1} \tilde{X}(t) + (ABA')^{-1} \tilde{U}(t), \quad t \geq 0,$$

where the Brownian motion $(ABA')^{-1} \tilde{X}$ has drift $(ABA')^{-1} \theta$ and covariance matrix $(ABA')^{-1} \Gamma (ABA')^{-1}$, and \tilde{U}_j can increase only when \tilde{Q}_j is zero, $j \in \mathbb{J}$. The




FIG 1. A linear network with two resources and three routes.

direction of reflection on the boundary face of \mathbb{R}_+^J is defined by the j^{th} column of the matrix $(ABA')^{-1}$. The initial distribution of \tilde{Q} is obtained by applying the linear transformation $(ABA')^{-1}$ to the distribution η . (For the formal definition of such an SRBM, where reflection directions are not in general parallel to coordinate directions, see [43].)

As an illustration of Theorem 5.2, consider a two resource linear network operating under a weighted proportional fair sharing policy ($\alpha = 1$). This network is depicted in Figure 1. It has two resources and three routes. Each resource has a route that passes only through that resource and then there is a route that passes through both resources.

The workload cone in this case is a wedge in \mathbb{R}_+^2 that has the following representation:

$$(51) \quad \mathcal{W}_1 = \left\{ \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \left(\frac{v_1}{\mu_1^2 \kappa_1} + \frac{v_3}{\mu_3^2 \kappa_3} \right) q_1 + \frac{v_3}{\mu_3^2 \kappa_3} q_2 \\ \left(\frac{v_2}{\mu_2^2 \kappa_2} + \frac{v_3}{\mu_3^2 \kappa_3} \right) q_2 + \frac{v_3}{\mu_3^2 \kappa_3} q_1 \end{bmatrix} : q \in \mathbb{R}_+^2 \right\}.$$

Let

$$\beta_1 = \left(1 + \frac{v_2 \mu_3^2 \kappa_3}{v_3 \mu_2^2 \kappa_2} \right) \quad \text{and} \quad \beta_2 = \left(1 + \frac{v_1 \mu_3^2 \kappa_3}{v_3 \mu_1^2 \kappa_1} \right).$$

It can be easily computed that the two boundary faces of the wedge \mathcal{W}_1 have the following expressions:

$$\mathcal{W}_1^1 = \{ w \in \mathbb{R}_+^2 : w_1 \geq 0, w_2 = \beta_1 w_1 \}$$

and

$$\mathcal{W}_1^2 = \{ w \in \mathbb{R}_+^2 : w_2 \geq 0, w_1 = \beta_2 w_2 \}$$




FIG 2. The workload cone \mathcal{W}_1 for a linear network with two resources and three routes is an infinite wedge in the positive quadrant of \mathbb{R}^2 . A finite portion of the wedge is shown above (between the two shaded regions). Under the lifting map Δ , points (w_1, w_2) on the boundary $w_2 = \beta_1 w_1$ of the wedge are mapped to points (n_1, n_2, n_3) where $n_1 = 0$ (and the corresponding $q \in \mathbb{R}_+^2$ has $q_1 = 0$); similarly, points (w_1, w_2) on the boundary $w_1 = \beta_2 w_2$ are mapped to points (n_1, n_2, n_3) where $n_2 = 0$ (and the corresponding $q \in \mathbb{R}_+^2$ has $q_2 = 0$).

The wedge is depicted in Figure 2. For a linear network, the local traffic condition holds automatically. Thus, the conclusion of Theorem 5.2 applies provided $\hat{W}^r(0)$ converges in distribution to a random variable with distribution η concentrated on \mathcal{W}_1 and $|\hat{N}^r(0) - \Delta(\hat{W}^r(0))| \rightarrow 0$ in probability as $r \rightarrow \infty$. For example, this holds if each of the systems indexed by r starts empty and η is the point mass at the origin in \mathbb{R}_+^2 . The limiting SRBM \hat{W} lives in the wedge \mathcal{W}_1 and is confined there by reflection (or pushing) at the boundary. Reflection occurs in the horizontal direction (corresponding to resource 1 underutilizing capacity) on the bounding face where $w_2 = \beta_1 w_1$. The interpretation of this is that although there is no work for resource 1 on route 1, there is work for this resource on route 3. However, congestion at resource 2, through the nature of the bandwidth sharing policy, is preventing resource 1 from working at its full capacity. Similarly, vertical reflection (corresponding to resource 2 underutilizing capacity) on the bounding face $w_1 = \beta_2 w_2$ is interpreted to mean that congestion at resource 1 is preventing resource 2 from working at its full capacity. Thus, the shape of the workload space indicates the entrainment of resources, whereby congestion at some resources may prevent other resources from working at their full capacity. We note that when $\kappa_3 \rightarrow \infty$ the upper boundary of \mathcal{W}_1 tends to the vertical axis and the lower boundary tends to the horizontal axis, hence the wedge expands to the whole quadrant, approaching the situation with full utilization of the resources.

5.4. *Product Form Stationary Distribution for Proportional Fair Sharing.* In this subsection we prove a result which shows that when $\alpha = 1$ and $\kappa_i = 1$ for all $i \in \mathbb{I}$ (proportional fair sharing), an SRBM \tilde{W} with the properties described in Theorem 5.2 has a product form invariant measure. (This is a result for SRBMs and so does not require the local traffic condition a priori.)

THEOREM 5.3. *Suppose that $\alpha = 1$ and $\kappa_i = 1$ for all $i \in \mathbb{I}$. Let π be the measure on \mathcal{W}_1 that is absolutely continuous with respect to Lebesgue measure with density given by*

$$(52) \quad p(w) = \exp(\mathbf{v} \cdot w), \quad w \in \mathcal{W}_1,$$

where

$$(53) \quad \mathbf{v} = 2\Gamma^{-1}\theta.$$

The product form measure π is an invariant measure for the SRBM having state space \mathcal{W}_1 , directions of reflection $\{\gamma^j : j \in \mathbb{J}\}$, drift θ , and covariance matrix Γ . This measure is integrable over \mathcal{W}_1 if and only if $\theta_j < 0$ for all $j \in \mathbb{J}$, and then after normalization it defines the unique stationary distribution for the SRBM.

PROOF. Sufficient conditions for a reflecting Brownian motion in a simple polyhedral domain to have a product form invariant measure were determined by Williams in [41], building on solutions of a related analytic problem obtained by Harrison and Williams in [17]. In these works, the covariance matrix for the process is the identity matrix. In order to apply these results, we need to perform a linear transformation to transform an SRBM with covariance matrix Γ into one whose covariance matrix is the identity matrix. We perform this transformation below; the computations are straightforward, though a little tedious, as we need to normalize the resulting directions of reflection to have inward normal components of unit length, to facilitate use of the results in [41]. Similar manipulations were carried out in the proof of Theorem 23 in [18] where the reflection directions had a special form.

Before introducing the transformation, we obtain an alternative representation for the simple convex polyhedron \mathcal{W}_1 . By (47),

$$\mathcal{W}_1 = \{ABA'q : q \in \mathbb{R}_+^{\mathbb{J}}\} = \{w \in \mathbb{R}^{\mathbb{J}} : (ABA')^{-1}w \in \mathbb{R}_+^{\mathbb{J}}\}.$$

Here, since $\kappa_i = 1$ for all i , we have that $B = M^{-1}\text{diag}(\mathbf{v})M^{-1}$ and by (46), $ABA' = \frac{1}{2}\Gamma$. Thus,

$$(54) \quad \mathcal{W}_1 = \{w \in \mathbb{R}^{\mathbb{J}} : \Gamma^{-1}w \in \mathbb{R}_+^{\mathbb{J}}\},$$

where we have used the fact that $q \in \mathbb{R}_+^{\mathbb{J}}$ if and only if $2q \in \mathbb{R}_+^{\mathbb{J}}$.

Now we define the linear transformation to be applied to the SRBM. Let Υ denote the diagonal matrix that has the same diagonal entries as Γ^{-1} . Let L be the rotation matrix whose rows are the orthonormal eigenvectors of the covariance matrix Γ and D be the corresponding diagonal matrix of eigenvalues such that $\Gamma = L'DL$ where $L' = L^{-1}$. Let \tilde{W} be an SRBM with state space \mathscr{W}_1 , drift θ , covariance matrix Γ and directions of reflection given by $\{\gamma^j : j \in \mathbb{J}\}$, with a decomposition as in (i) of Definition 5.2. Let $V = D^{-\frac{1}{2}}L$ and define $\tilde{Z} = V\tilde{W}$. Then \tilde{Z} is an SRBM in the simple convex polyhedron

$$\begin{aligned} \mathscr{Z}_1 &= \{Vw : w \in \mathscr{W}_1\} = \{z \in \mathbb{R}^{\mathbb{J}} : V^{-1}z \in \mathscr{W}_1\} \\ &= \{z \in \mathbb{R}^{\mathbb{J}} : \Gamma^{-1}V^{-1}z \in \mathbb{R}_+^{\mathbb{J}}\} \\ &= \{z \in \mathbb{R}^{\mathbb{J}} : \tilde{n}^j \cdot z \geq 0 \text{ for all } j \in \mathbb{J}\} \end{aligned}$$

where for each $j \in \mathbb{J}$, \tilde{n}^j is given by the j^{th} row of the matrix

$$(55) \quad \Theta = \Upsilon^{-\frac{1}{2}}\Gamma^{-1}V^{-1} = \Upsilon^{-\frac{1}{2}}\Gamma^{-1}L'D^{\frac{1}{2}},$$

and \tilde{n}^j is the unit inward normal to the j^{th} face of \mathscr{Z}_1 . (The matrix $\Upsilon^{-\frac{1}{2}}$ is used to normalize the \tilde{n}^j to be of unit length.) The process \tilde{Z} inherits the following decomposition from \tilde{W} :

$$(56) \quad \tilde{Z}(t) = \tilde{Z}(0) + V\tilde{X}(t) + R\tilde{Y}(t), \quad t \geq 0, \quad \text{a.s.}$$

where $R = VH^{-1} = D^{-\frac{1}{2}}LH^{-1}$, $H = \Upsilon^{\frac{1}{2}}$ is the diagonal matrix with the same diagonal entries as ΘV , and $\tilde{Y} = H\tilde{U}$ is a continuous, non-decreasing process that starts from zero and can increase only when \tilde{Z} is on the j^{th} face of \mathscr{Z}_1 . The columns of the matrix R are the directions of reflection for \tilde{Z} on the faces of \mathscr{Z}_1 , normalized so that the inward normal component of each direction of reflection has unit length. (The matrix H^{-1} is used to achieve the normalization.) The matrix R has the form

$$(57) \quad R = \Theta' + \Xi'$$

where Θ is the matrix specified in (55) (whose rows are the inward unit normals to the faces of \mathscr{Z}_1) and Ξ is the matrix whose rows consist of the components of the (normalized) directions of reflection that are tangent to each of the faces of \mathscr{Z}_1 . In particular, the diagonal entries of $\Theta\Xi'$ are all zero. In this context, the sufficient condition given in [41] for existence of a product form invariant measure for \tilde{Z} is the so-called ‘‘skew symmetry condition’’:

$$(58) \quad \Theta\Xi' + \Xi\Theta' = 0.$$

The vector form of (58) is

$$(59) \quad \tilde{n}^i \cdot \tilde{s}^j + \tilde{s}^i \cdot \tilde{n}^j = 0 \quad \text{for all } i, j \in \mathbb{J},$$

where for each $i \in \mathbb{J}$, \tilde{n}^i is the inward unit normal to the i^{th} face of \mathcal{Z}_1 and \tilde{s}^i is the tangential component of the (normalized) direction of reflection on that face. A geometric interpretation of the skew symmetry condition is provided in [17]. We refer the reader to that paper, especially Section 9, for a full discussion. Briefly, for a simple convex polyhedron as we have here, the skew symmetry condition can be shown to be equivalent to a local condition. This local condition can be stated in words as requiring that near the intersection of any two distinct faces of \mathcal{Z}_1 , the component of the reflection vector on one face that is directed towards the intersection of the faces is balanced on the other face by the component of the direction of reflection for that second face which is of the same magnitude as the component on the first face, but it is directed away from the intersection of the faces. Indeed, under the skew symmetry condition, with probability one, the SRBM \tilde{Z} will not hit the intersection of any two faces when started away from that set, as is proved in [41].

The left hand side of condition (58) is equal to

$$\begin{aligned}
\Theta(R - \Theta') + (R' - \Theta)\Theta' &= -2\Theta\Theta' + \Theta R + R'\Theta' \\
&= -2\Upsilon^{-\frac{1}{2}}\Gamma^{-1}L'D^{\frac{1}{2}}D^{\frac{1}{2}}L\Gamma^{-1}\Upsilon^{-\frac{1}{2}} \\
&\quad + \Upsilon^{-\frac{1}{2}}\Gamma^{-1}L'D^{\frac{1}{2}}D^{-\frac{1}{2}}LH^{-1} \\
&\quad + H^{-1}L'D^{-\frac{1}{2}}D^{\frac{1}{2}}L\Gamma^{-1}\Upsilon^{-\frac{1}{2}} \\
&= -2\Upsilon^{-\frac{1}{2}}\Gamma^{-1}\Upsilon^{-\frac{1}{2}} \\
&\quad + \Upsilon^{-\frac{1}{2}}\Gamma^{-1}H^{-1} + H^{-1}\Gamma^{-1}\Upsilon^{-\frac{1}{2}} \\
&= 0,
\end{aligned}$$

where we have used the facts that $\Gamma = L'DL$, $L'L = I$ and $H = \Upsilon^{\frac{1}{2}}$.

Thus, the skew symmetry condition (58) holds and it follows from Theorem 1.2 of [41] that the SRBM \tilde{Z} has a product form invariant measure with a density relative to Lebesgue measure that is proportional to $\exp(\beta \cdot z)$, $z \in \mathcal{Z}_1$, where $\beta = 2(I - \Theta^{-1}\Xi)^{-1}V\theta = 2V\theta$. Note for this that

$$\begin{aligned}
\Theta^{-1}\Xi &= D^{-\frac{1}{2}}L\Gamma\Upsilon^{\frac{1}{2}}(H^{-1}L'D^{-\frac{1}{2}} - \Upsilon^{-\frac{1}{2}}\Gamma^{-1}L'D^{\frac{1}{2}}) \\
&= D^{-\frac{1}{2}}L\Gamma\Upsilon^{\frac{1}{2}}(\Upsilon^{-\frac{1}{2}}L'D^{-\frac{1}{2}} - \Upsilon^{-\frac{1}{2}}L'D^{-1}LL'D^{\frac{1}{2}}) \\
&= D^{-\frac{1}{2}}L\Gamma\Upsilon^{\frac{1}{2}}(\Upsilon^{-\frac{1}{2}}L'D^{-\frac{1}{2}} - \Upsilon^{-\frac{1}{2}}L'D^{-\frac{1}{2}}) \\
&= 0,
\end{aligned}$$

where we have used the facts that $\Gamma^{-1} = L'D^{-1}L$ and $L'L = I$. Furthermore, by Corollary 1.1 of [41], if the exponential density is integrable over \mathcal{Z}_1 , then after normalization it yields the unique stationary distribution for \tilde{Z} .

By inverting the linear transformation V , we can transform this result back to one for \tilde{W} . Noting that for $z = Vw$,

$$(60) \quad \beta \cdot z = \beta \cdot Vw = 2\theta'V'Vw = 2\theta'\Gamma^{-1}w,$$

we conclude that (52) is an invariant density for the original SRBM \tilde{W} and if this is integrable over \mathscr{W}_1 , then after normalization it yields the unique stationary distribution for \tilde{W} . Using the representation (54) for \mathscr{W}_1 , we see that this density will be integrable over \mathscr{W}_1 if and only if $\exp(\theta \cdot q)$, $q \in \mathbb{R}_+^{\mathbb{J}}$ is integrable over $\mathbb{R}_+^{\mathbb{J}}$, which occurs if and only if $\theta_j < 0$ for each $j \in \mathbb{J}$. \square

REMARK 5.4. *The product form invariant measure of Theorem 5.3 is remarkable, yet the proof of the result gives little insight into why the reflection directions and covariance matrix in the particular case $\alpha = 1$, $\kappa_i = 1$ for $i = 1, \dots, \mathbb{I}$, should allow the result to hold. We simply note here that the authors first suspected that a product form result might be found after observing that it is possible to describe a network of queues with a product form stationary distribution whose conjectured Brownian model approximation has the same directions of reflection and covariance matrix as the Brownian model approximation under study in this paper, provided $\alpha = 1$ and $\kappa_i = 1$ for $i = 1, \dots, \mathbb{I}$, i.e., the case of proportional fair sharing. Earlier connections between product form queueing networks and proportional fairness have been explored by [4, 34], and the relationship between these several connections seems a rich area for further study.*

The product form of the density (52) does not imply that, when $\theta_j < 0$ for all $j \in \mathbb{J}$, the components of the SRBM \tilde{W} are independent under the stationary distribution for the SRBM, since in general the cone \mathscr{W}_1 is not an orthant. Independence can, however, be deduced for the components of the SRBM \tilde{Q} of dual variables.

COROLLARY 5.1. *Suppose that the assumptions of Theorem 5.2 hold, that $\kappa_i = 1$ for all $i \in \mathbb{I}$, and that $\theta_j < 0$ for all $j \in \mathbb{J}$. Let (\tilde{W}, \tilde{N}) be the process identified in Theorem 5.2. Then the SRBM $\tilde{Q} = 2\Gamma^{-1}\tilde{W}$ of dual variables has a unique stationary distribution, and this distribution has a density relative to Lebesgue measure that is proportional to $\exp(\theta \cdot q)$, $q \in \mathbb{R}_+^{\mathbb{J}}$. Under this stationary distribution the components of \tilde{Q} are independent and \tilde{Q}_j is exponentially distributed with parameter $-\theta_j$ for each $j \in \mathbb{J}$.*

PROOF. This result is immediate from Theorem 5.3 upon applying the linear transformation $(ABA')^{-1} = 2\Gamma^{-1}$ to transform \tilde{W} into \tilde{Q} . \square

Under the assumptions of Corollary 5.1, from Theorem 5.2, (48) and the definition of \tilde{Q} we have that

$$(61) \quad \tilde{N} = \text{diag}(\rho)A'\tilde{Q},$$

and it follows that the stationary distribution of \tilde{N} can be expressed as a linear combination of independent exponential random variables. Thus resource j has associated with it a dual random variable \tilde{Q}_j^s , for $j \in \mathbb{J}$ (here the superscript of s signals that the random variable is associated with the stationary distribution); these dual variables are independent and exponentially distributed with parameters $-\theta_j$ for $j \in \mathbb{J}$; and under its stationary distribution the i^{th} component of \tilde{N} is proportional to the sum of the dual variables associated with the resources used by route i . This suggests the following simple approximation for the stationary distribution of the *unscaled* network, that is the flow level model of Section 2. The stationary approximation is

$$(62) \quad N_i^s \approx \rho_i \sum_{j \in \mathbb{J}} Q_j^s A_{ji}$$

where Q_j^s , $j \in \mathbb{J}$, are independent and Q_j^s is exponentially distributed with parameter $C_j - \sum_{i \in \mathbb{I}} A_{ji} \rho_i$. We want to emphasize that (62) above is merely a formal approximation involving various implicit assumptions such as existence of a stationary distribution for N and formal unraveling of the heavy traffic scaling and limit procedure. In particular, this involves an interchange of limits and the approximation in (62) involves errors that may be substantial, e.g., of order r . However, the following observation of Massoulié and Roberts [33], which yields an exact stationary distribution for an unscaled linear network, suggests that there is cause for optimism regarding the approximation (62).

EXAMPLE 5.1. *Consider a linear network with \mathbf{J} resources where the set of resources is $\mathbb{J} = \{1, 2, \dots, \mathbf{J}\}$. Let the set of routes be labelled $\mathbb{I} = \{0, 1, 2, \dots, \mathbf{J}\}$ where we use the symbol $i = 0$ for the route $\{1, 2, \dots, \mathbf{J}\}$ through every resource, and for $i = 1, 2, \dots, \mathbf{J}$ we use the symbol i for the route $\{i\}$ through the single resource i . The local traffic Assumption 5.1 thus holds, and we assume $\alpha = 1$ and $\kappa_i = 1$ for all $i \in \mathbb{I}$. Suppose that $C_j = 1$, $j = 1, 2, \dots, \mathbf{J}$, and that $\rho_0 + \rho_j < 1$, $j = 1, 2, \dots, \mathbf{J}$. Then the stationary distribution for $(N_0, N_1, \dots, N_{\mathbf{J}})$ is given by [33]:*

$$(63) \quad \pi(n_0, n_1, \dots, n_{\mathbf{J}}) = \frac{\prod_{j=1}^{\mathbf{J}} (1 - \rho_0 - \rho_j)}{(1 - \rho_0)^{\mathbf{J}-1}} \binom{\sum_{i=0}^{\mathbf{J}} n_i}{n_0} \prod_{i=0}^{\mathbf{J}} \rho_i^{n_i},$$

where $n_0, \dots, n_{\mathbf{J}}$ each run through the non-negative integers. Summing this formula over n_0 , and using the negative binomial expansion, gives the marginal distribution for $(N_1, N_2, \dots, N_{\mathbf{J}})$ as

$$(64) \quad \pi(n_1, n_2, \dots, n_{\mathbf{J}}) = \prod_{i=1}^{\mathbf{J}} \frac{(1 - \rho_0 - \rho_i)}{(1 - \rho_0)} \left(\frac{\rho_i}{1 - \rho_0} \right)^{n_i}.$$

Thus, under the stationary distribution, $N_1, N_2, \dots, N_{\mathbf{J}}$ are independent, and N_i is geometrically distributed with mean $\rho_i / (1 - \rho_0 - \rho_i)$. This accords remarkably well with the approximation (62), under which N_i would be approximated by an exponentially distributed random variable with the same mean.

For this particular example, as observed by [33], the Markov chain N is reversible and consequently the stationary distribution π for N is insensitive to the document size distributions, depending only on their means. (For a description of general system structures to which such insensitivity results apply, see [37] and references therein.)

As a complement to the insensitivity result mentioned in the above example, we note that the stationary distribution result of Corollary 5.1 can be extended to the situation where the document sizes are finite mixtures of exponentials. Indeed, such a flow level model may be realized by collapsing an extended (exponential) model. The extended model is obtained by splitting each route in the original model into finitely many ‘‘copies’’, so that each copy uses the same set of resources and has the same weights as the original route but where each copy has its own Poisson arrival process and distinct exponential document size distribution. The relative arrival rates for the copies determine the proportions for the mixture of exponential distributions associated with the original route. More precisely, consider an extended (exponential) flow level model in which, for each $i = 1, \dots, \mathbf{I}$, routes $i(1), i(2), \dots, i(K_i)$ identify finitely many identical routes (subsets of the resources in \mathbb{J}) whose associated weights are also identical. In the following, parameters and processes associated with this extended model will have a superscript of \dagger appended. In particular, N^\dagger denotes the flow count process.

Under the assumptions of Theorem 5.2 (heavy traffic, $\alpha = 1$, local traffic condition and initial state space collapse) for the extended flow level model, there is an SRBM approximation \tilde{W}^\dagger for the workload and an attendant approximation $\tilde{N}^\dagger = \Delta(\tilde{W}^\dagger)$ for the flow count process. Furthermore, there is an SRBM process \tilde{Q}^\dagger of dual random variables. Under the additional assumptions of Corollary 5.1 (all of the weights are one and $\theta_j^\dagger < 0$ for each $j \in \mathbb{J}$), the stationary distribution of \tilde{Q}^\dagger is such that the components are independent and \tilde{Q}_j^\dagger is exponentially distributed with parameter $-\theta_j^\dagger$ for each $j \in \mathbb{J}$. Thus, under the assumptions of Corollary 5.1, the stationary distribution of

$$(65) \quad \tilde{N}^\dagger = \text{diag}(\rho^\dagger)(A^\dagger)' \tilde{Q}^\dagger,$$

can be expressed as a linear combination of exponential random variables.

For the extended model, the Poisson arrival rate for route $i(k)$ is $v_{i(k)}^\dagger$ and the exponential service time parameter for this route is $\mu_{i(k)}^\dagger$. Then, the \mathbf{I} -dimensional

collapsed process N defined by

$$(66) \quad N_i = \sum_{k=1}^{K_i} N_{i(k)}^\dagger, \quad i = 1, \dots, \mathbf{I},$$

has the same distribution as the flow count process in a flow level model with \mathbf{I} routes where the Poisson arrival rate for route $i = 1, \dots, \mathbf{I}$ is

$$(67) \quad v_i = \sum_{k=1}^{K_i} v_{i(k)}^\dagger$$

and the document size distribution for route $i = 1, \dots, \mathbf{I}$ is a mixture of K_i exponentials with parameters $\mu_{i(k)}^\dagger$ and proportions $v_{i(k)}^\dagger/v_i$ for $k = 1, \dots, K_i$. The mean of this document size distribution is $1/\mu_i$ where μ_i is defined by

$$(68) \quad \frac{1}{\mu_i} = \sum_{k=1}^{K_i} \frac{v_{i(k)}^\dagger}{v_i} \frac{1}{\mu_{i(k)}^\dagger}$$

for $i = 1, \dots, \mathbf{I}$. Let $\rho_i = v_i/\mu_i$, for $i = 1, \dots, \mathbf{I}$.

Now consider the exponential analogue of the collapsed flow level model where the finite mixture of exponential distributions for route i is replaced by a single exponential distribution with parameter μ_i for $i = 1, \dots, \mathbf{I}$. The nominal load placed on resource j in the extended model is $(A^\dagger \rho^\dagger)_j$ and this is the same as the nominal load $(A\rho)_j$ placed on j in the exponential analogue. It follows that a sequence of extended flow level models satisfying the assumptions of Corollary 5.1, where the limiting value in (12) of the heavy traffic Assumption 3.1 is denoted by θ^\dagger , will have a parallel sequence of exponential analogues whose limiting value θ in (12) will be precisely the same as θ^\dagger . Consequently, the stationary distribution of \tilde{Q}^\dagger is the same as that for the process \tilde{Q} of dual random variables associated with the sequence of exponential analogues.

Then, combining (65) and (66), we obtain the following as an approximation for the collapsed flow count process:

$$\tilde{N}_i = \sum_{k=1}^{K_i} \rho_{i(k)}^\dagger \sum_{j \in \mathbb{J}} A_{ji(k)}^\dagger \tilde{Q}_j^\dagger = \rho_i \sum_{j \in \mathbb{J}} A_{ji} \tilde{Q}_j^\dagger, \quad \text{for } i = 1, \dots, \mathbf{I},$$

that is, $\tilde{N} = \text{diag}(\rho)A'\tilde{Q}^\dagger$. Since the stationary distributions for \tilde{Q}^\dagger and \tilde{Q} are the same, it follows that the stationary distribution for \tilde{N} is the same whether one considers the collapsed flow level model (where document sizes are distributed as finite mixtures of exponentials) or the associated exponential analogue (where document sizes are distributed as exponentials). The formal approximation (62) for the unscaled network is similarly the same for both. Note, however, that the stochastic

processes \tilde{W}^\dagger , \tilde{N} , and \tilde{Q}^\dagger , and in particular their covariance matrices, are in general different from those associated with the exponential analogue.

It is natural to conjecture an extension of the above insensitivity results to general document size distributions. However, even an extension to phase-type distributions would require generalization of the diffusion approximation results to flow level models with feedback, and treatment of more general distributions would appear to require a significantly more elaborate stochastic model (cf. [14]) to keep track of residual document sizes.

5.5. Multi-Path Routing. In this subsection we describe a generalization of the earlier model that allows multi-path routing. In our initial description of the model we shall use a different notation for the sets of routes and resources: this will allow our eventual results to be expressed in a notation consistent with that used elsewhere in the paper.

Suppose that now we interpret $i \in \mathbb{I}$ as a source-destination pair, and let \mathbb{K} be the set of routes. Let $\mathbf{I} = |\mathbb{I}|$, and let $\mathbf{K} = |\mathbb{K}|$. We suppose that \mathbb{K} is partitioned into \mathbf{I} non-empty subsets, each associated with a single source-destination pair. Let $H_{ik} = 1$ if route $k \in \mathbb{K}$ is associated with source-destination pair $i \in \mathbb{I}$, and let $H_{ik} = 0$ otherwise. Thus H is a $\mathbf{I} \times \mathbf{K}$ matrix containing only zeros and ones, and $\sum_{i \in \mathbb{I}} H_{ik} = 1$ for each $k \in \mathbb{K}$.

There remain finitely many resources, but now labeled by $l \in \mathbb{L}$, with capacities $(\bar{C}_l : l \in \mathbb{L})$ that are all strictly positive and finite. A route k is a non-empty subset of \mathbb{L} (interpreted as the set of resources used by route k). We assume that \mathbb{L} and \mathbb{K} are both non-empty and finite. Let $\mathbf{L} = |\mathbb{L}|$, the total number of resources. Let \bar{A} be the $\mathbf{L} \times \mathbf{K}$ matrix containing only zeros and ones, defined such that $\bar{A}_{lk} = 1$ if resource l is used by route k , and $\bar{A}_{lk} = 0$ otherwise. Our assumption that each route k identifies a *non-empty* subset of \mathbb{L} implies that no column of \bar{A} is identically zero.

It is assumed that a new document arrives to source-destination pair i at each jump time of a Poisson process that has rate parameter $\nu_i > 0$ and that each such document has an exponentially distributed size with mean $1/\mu_i$ where $\mu_i \in (0, \infty)$. These document sizes are assumed to be independent of one another and to be independent of all arrival times of documents. Let $\rho_i = \nu_i/\mu_i, i \in \mathbb{I}$.

Given a fixed parameter $\alpha \in (0, \infty)$ and strictly positive weights $(\kappa_i : i \in \mathbb{I})$, if $N_i(t)$ denotes the (random) number of documents being transferred between source-destination pair i at time t for each $i \in \mathbb{I}$ and $N(t) = (N_i(t) : i \in \mathbb{I})$, then the bandwidth allocated to source-destination pair i at time t is given by $\Lambda_i(N(t))$ and this bandwidth is shared equally amongst the $N_i(t)$ transfers in progress between source-destination pair i . The function $\Lambda = (\Lambda_i : i \in \mathbb{I}) : \mathbb{R}_+^{\mathbf{I}} \rightarrow \mathbb{R}_+^{\mathbf{I}}$ is defined such that for each $n \in \mathbb{R}_+^{\mathbf{I}}$, $\Lambda_i(n) = 0$ for $i \in \mathbb{I}_0(n) \equiv \{m \in \mathbb{I} : n_m = 0\}$, and when

$\mathbb{I}_+(n) \equiv \{m \in \mathbb{I} : n_m > 0\}$ is non-empty, $\Lambda^+(n) \equiv (\Lambda_i(n) : i \in \mathbb{I}_+(n))$ is the unique value of $\Lambda^+ = (\Lambda_i : i \in \mathbb{I}_+(n))$ that solves the optimization problem:

$$(69) \quad \begin{array}{ll} \text{maximize} & G_n(\Lambda^+) \\ \text{subject to} & \sum \bar{A}_{lk} y_k \leq \bar{C}_l, \quad l \in \mathbb{L}, \\ & \sum_k H_{ik} y_k = \Lambda_i, \quad i \in \mathbb{I}_+(n), \\ \text{over} & y_k \geq 0, \quad k \in \mathbb{K}, \quad \Lambda_i \geq 0, \quad i \in \mathbb{I}_+(n), \end{array}$$

where $G_n(\Lambda^+)$ is again given by the definition (2).

Without loss, we may assume that in the solution of (69), $y_k = 0$ for those k such that $H_{ik} = 0$ for all $i \in \mathbb{I}_+(n)$. With this convention, allocations $y_k, k \in \mathbb{K}$ associated with the unique optimal value $\Lambda^+(n)$ can be interpreted as bandwidth allocations for the routes which sum up to give the bandwidth allocations to the source-destination pairs.

The above optimization problem reduces to the earlier problem (1) in the case where $\mathbf{I} = \mathbf{K}$ and H is the identity matrix, that is, the case of a single route for each source-destination pair. More generally, the following proposition allows us to reduce to the optimization problem (1) even in the multi-path case. This observation was previously made in [24], Section 3.3.

Let

$$\mathcal{Y} = \{y \in \mathbb{R}_+^{\mathbf{K}} : \bar{A}y \leq \bar{C}\}.$$

PROPOSITION 5.1. *There exists a representation*

$$(70) \quad H\mathcal{Y} = \{(\Lambda_i, i \in \mathbb{I}) : \Lambda_i \geq 0, i \in \mathbb{I}, \sum_{i \in \mathbb{I}} A_{ji} \Lambda_i \leq C_j, j \in \mathbb{J}\}$$

where \mathbb{J} , A and C can be chosen so that C has positive elements and A has non-negative elements and no column of A is identically zero.

PROOF. The set \mathcal{Y} is the intersection of the half-spaces $\{y \in \mathbb{R}^{\mathbf{K}} : (\bar{A}y)_l \leq \bar{C}_l\}$, $l \in \mathbb{L}$, and the non-negative orthant, $\mathbb{R}_+^{\mathbf{K}}$. Since no column of \bar{A} is identically zero, the set \mathcal{Y} is bounded and is thus the convex hull of a finite number of extreme points. Hence $H\mathcal{Y}$ is the convex hull of a finite number of extreme points, or, equivalently, the bounded intersection of a finite set of half-spaces.

Next we explore further the geometry of $H\mathcal{Y}$. First, $H\mathcal{Y} \subset \mathbb{R}_+^{\mathbf{I}}$, since the elements of H are non-negative; also $0 \in H\mathcal{Y}$, since $0 \in \mathcal{Y}$. Indeed for δ positive and small enough, $(\Lambda_i = \delta_i, i \in \mathbb{I}) \in H\mathcal{Y}$ for all $0 \leq \delta_i < \delta$, $i \in \mathbb{I}$, since for ε positive and small enough, $(y_k = \varepsilon_k, k \in \mathbb{K}) \in \mathcal{Y}$ for all $0 \leq \varepsilon_k < \varepsilon$, $k \in \mathbb{K}$, and no row of H is identically zero. Thus $H\mathcal{Y}$ is bounded by the hyperplanes bounding the non-negative orthant, plus finitely many other hyperplanes, none of which contains

the origin. Thus there exists a representation of the form (70), for some choice of \mathbb{J} , A and C . Since $0 \in H\mathcal{Y}$, the elements of C are non-negative. Choose a minimal representation, where the set \mathbb{J} is of minimal cardinality; then the intersection of the hyperplane

$$(71) \quad \{(\Lambda_i, i \in \mathbb{I}) : \sum_{i \in \mathbb{I}} A_{ji} \Lambda_i = C_j\}$$

with $H\mathcal{Y}$ must have a non-empty interior relative to the hyperplane, for each $j \in \mathbb{J}$.

Next we show that if $\Lambda \in H\mathcal{Y}$, and $0 \leq \Lambda' \leq \Lambda$, then $\Lambda' \in H\mathcal{Y}$. Let $y \in \mathcal{Y}$ be such that $\Lambda = Hy$, and for each $k \in \mathbb{K}$ let $i(k)$ be the unique index i such that $H_{ik} = 1$. Let $y'_k = (\Lambda'_{i(k)} / \Lambda_{i(k)}) y_k$ for $k \in \mathbb{K}$. Then $y' \in \mathcal{Y}$ and $Hy' = \Lambda'$, and so $\Lambda' \in H\mathcal{Y}$.

We have seen that the intersection of the hyperplane (71) with $H\mathcal{Y}$ has a non-empty interior relative to the hyperplane; choose a point Λ in this relative interior. Then $\Lambda_i > 0$ for $i \in \mathbb{I}$. Fix $i \in \mathbb{I}$ and choose Λ' so that $0 < \Lambda'_i < \Lambda_i$, with $\Lambda'_j = \Lambda_j$ for $j \in \mathbb{I} \setminus \{i\}$. Then $0 \leq \Lambda' \leq \Lambda$, and so $\Lambda' \in H\mathcal{Y}$. Hence the hyperplane (71) must have $A_{ji} \geq 0$.

Finally note that $C_j > 0$, for $j \in \mathbb{J}$, since the hyperplane (71) does not contain the origin, and no column of A is identically zero, since $H\mathcal{Y}$ is bounded. \square

The representation (70) of $H\mathcal{Y}$ allows us to elide the variables y from the optimization definition of $\Lambda^+(n)$, and to deduce that the unique solution $(\Lambda_i : i \in \mathbb{I}_+(n))$ to the optimization problem (69) is also the unique solution to the optimization problem (1). With this generalization to multi-path routing, the matrix A may no longer only contain zeros and ones. However, the proofs of Theorem 5.2 and Corollary 5.1 still go through in this enhanced generality provided the local traffic condition is satisfied.

As an illustration consider the network depicted in Figure 3, operating under the proportional fair sharing policy ($\alpha = 1$, $\kappa_i = 1$ for all i), where the labelled resource capacities satisfy $\bar{C}_3 < \bar{C}_1, \bar{C}_2$. It has three source-destination pairs, four resources and five routes. Then A and C can be chosen as

$$A = \begin{pmatrix} 1 & 1 & 0 \\ \frac{1}{2} & 0 & 1 \end{pmatrix} \quad C = \begin{pmatrix} \bar{C}_1 + \bar{C}_2 \\ \bar{C}_3 \end{pmatrix}.$$

The matrices A and C may be viewed as expressing the *generalized cut constraints*

$$(72) \quad \rho_1 + \rho_2 \leq \bar{C}_1 + \bar{C}_2$$

$$(73) \quad \frac{1}{2} \rho_1 + \rho_3 \leq \bar{C}_3$$

apparent from Figure 3. In the first of these constraints, the capacities of the first and second resources are pooled. Resource pooling subject to generalized cut constraints is a common phenomenon in stochastic processing networks with routing [29]. Theorem 3.1 of [29] provides an alternative form of the representation (70), and provides references to algorithms to calculate matrices A and C .

$$(74) \quad \mathcal{W}_\alpha = \left\{ \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} = \begin{bmatrix} \frac{v_1}{\mu_1^2 \kappa_1^\alpha} q_1^{\frac{1}{\alpha}} + \frac{v_3}{\mu_3^2 \kappa_3^\alpha} (q_1 + q_2)^{\frac{1}{\alpha}} \\ \frac{v_2}{\mu_2^2 \kappa_2^\alpha} q_2^{\frac{1}{\alpha}} + \frac{v_3}{\mu_3^2 \kappa_3^\alpha} (q_1 + q_2)^{\frac{1}{\alpha}} \end{bmatrix} : q \in \mathbb{R}_+^2 \right\}.$$

The two boundary faces of this wedge have the following expressions:

$$\mathcal{W}_\alpha^1 = \left\{ \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} : w_1 \geq 0, w_2 = \left(1 + \frac{\mu_2^{-2} v_2}{\mu_3^{-2} v_3} \left(\frac{\kappa_3}{\kappa_2} \right)^{1/\alpha} \right) w_1 \right\},$$

and

$$\mathcal{W}_\alpha^2 = \left\{ \begin{bmatrix} w_1 \\ w_2 \end{bmatrix} : w_2 \geq 0, w_1 = \left(1 + \frac{\mu_1^{-2} v_1}{\mu_3^{-2} v_3} \left(\frac{\kappa_3}{\kappa_1} \right)^{1/\alpha} \right) w_2 \right\}.$$

In this case, since the two-dimensional wedge \mathcal{W}_α is still a polyhedral cone, the proof of Theorem 5.2 can be easily modified to apply for $\alpha \neq 1$, with \mathcal{W}_α used in place of \mathcal{W}_1 there. When the weights $\kappa_i, i = 1, 2, 3$, are all equal, the wedge \mathcal{W}_α does not depend on α . However, when the κ_i 's are not all equal, the wedge \mathcal{W}_α does depend on α . However, as α goes to infinity, the quotients involving the weights κ_i in the expressions for $\mathcal{W}_\alpha^j, j = 1, 2$, tend to one and there is a limiting wedge which is the same as that obtained when the κ_i are all equal. On the other hand, as α tends to zero, the limit of the upper boundary depends on whether $\kappa_3 > \kappa_2$ (tends to the vertical axis) or $\kappa_3 < \kappa_2$ (tends to the 45 degree ray from the origin) and similarly for the lower boundary: if $\kappa_3 > \kappa_1$, it tends to the horizontal axis and if $\kappa_3 < \kappa_1$, it tends to the 45 degree ray from the origin. Hence, as α tends to zero, if $\kappa_3 > \max\{\kappa_1, \kappa_2\}$, the wedge expands to the whole quadrant, and if $\kappa_3 < \min\{\kappa_1, \kappa_2\}$, the wedge contracts to the 45 degree ray from the origin. Note that even when the wedge expands to the whole quadrant, the components of the diffusion workload do not become independent as the covariance matrix Γ , which does not depend on κ or α , is not diagonal.

In general, for $\alpha \neq 1$ with higher dimensional workloads, i.e., $\mathbf{J} > 2$, the shape of the workload cone \mathcal{W}_α depends on α , as well as on A, v, μ and κ . This relationship appears to be quite complicated. We are investigating Conjecture 5.1 in this case. At this time we only have some partial results. The main difficulty in establishing the validity of the conjecture concerns proving C -tightness of $\{\hat{W}^r\}$ and establishing uniqueness in law for any weak limit point. To illustrate some of the difficulties involved, we consider a linear network with three resources and four routes. Although this is a particular example, it exemplifies the challenges presented by the geometric data in establishing existence, uniqueness and an invariance principle for the SRBM. For this example, each of the three resources has a route that only goes




FIG 4. A portion of the workload cone \mathcal{W}_1 is shown for a linear network with three resources and four routes with $\alpha = 1$ and $v_i = \mu_i = \kappa_i = 1$ for all $i \in \mathbb{I}$.




FIG 5. A cross-section of the workload cone \mathcal{W}_1 depicted in Figure 4 taken at $w_3 = 1$.
imsart-aap ver. 2006/03/07 file: revcongestion.tex date: August 4, 2008




FIG 6. A portion of the workload cone \mathcal{W}_2 is shown for a linear network with three resources and four routes with $\alpha = 2$ and $v_i = \mu_i = \kappa_i = 1$ for all $i \in \mathbb{I}$.




FIG 7. A cross-section of the workload cone \mathcal{W}_2 depicted in Figure 6 taken at $w_3 = 1$.




FIG 8. A portion of the workload cone $\mathcal{W}_{0.5}$ is shown for a linear network with three resources and four routes with $\alpha = 0.5$ and $v_i = \mu_i = \kappa_i = 1$ for all $i \in \mathbb{I}$.




FIG 9. A cross-section of the workload cone $\mathcal{W}_{0.5}$ depicted in Figure 8 taken at $w_3 = 1$.

through that resource and there is one additional route that goes through all three resources. Here the workload cone is contained in the three dimensional positive orthant. When $v_i = \mu_i = \kappa_i = 1$, for all $i \in \mathbb{I}$, we depict the associated workload cone, and one of its cross-sections, for $\alpha = 2$, $\alpha = 1$ and $\alpha = 0.5$ in Figures 4–9.

For the case depicted in Figures 4–5, $\alpha = 1$ and Theorem 5.2 applies to justify the diffusion approximation. The proof of this theorem uses the uniqueness of the diffusion process [10] and the invariance principle developed in [20]. For the case depicted in Figure 6–7, $\alpha = 2$ and the workload cone has boundary faces that curve outwards. Away from the origin, the boundary and directions of reflection for the proposed diffusion approximation locally satisfy conditions required by the invariance principle given in [20]. However, the workload cone has a “singular point” at the origin where the conditions in [20] fail to be satisfied. However, since this is an isolated point, the invariance principle in [20] and the proof of uniqueness starting in [10] can be adapted to establish the SRBM diffusion approximation in this case. For higher dimensional analogues of this case, we anticipate that a valid diffusion approximation can be established. However, as the dimension increases, it becomes more difficult to compute the inward normals to all boundary faces, and as yet we do not have a systematic way to verify geometric sufficient conditions for uniqueness in law and a valid invariance principle. For the case depicted in Figures 8–9, where $\alpha = 0.5$, the situation is more complex. Here the workload cone has boundary faces that curve inward and any two faces meet in a cusp, i.e., the inward normals to any two boundary faces point towards one another at the intersection of the faces. At present we do not have an existence and uniqueness result, nor an invariance principle, to treat this case because of the singular geometry at intersections of faces. In particular, it has not been established that the diffusion process can escape from the tip of a cusp or that there is uniqueness in law for a diffusion process starting there.

In a recent work, Shah and Wischik [38] have proved multiplicative state space collapse for a class of “switched” networks operating under a family of scheduling policies related to the maximum weight algorithm introduced by Tassiulas and Ephremides [40]. Based on this, Shah and Wischik have conjectured natural diffusion approximations for the workload processes in these models. These diffusions are SRBMs living in cones with piecewise smooth boundaries. While these cones have some characteristics in common with those found for the bandwidth sharing model considered here, there are also some new features. Depending on a parameter associated with the family of scheduling policies, the cones of [38] include non-simple convex polyhedrons as well as cones with curved boundaries where boundary faces can meet smoothly. The validity of these conjectured diffusion approximations for input queued packet switches operating under a maximum weight algorithm is being explored in [21].

6. PROOF OF MULTIPLICATIVE STATE SPACE COLLAPSE. In this section we prove the multiplicative state space collapse result, Theorem 5.1. Our proof follows a general line of argument pioneered by Bramson in [5], where open multiclass queueing networks operating under certain head-of-the-line (HL) service disciplines are treated. However, there are some differences from [5]. Here we have the more general structure of a stochastic processing network with simultaneous resource possession, and our service discipline is not work conserving. On the other hand, we have exponential interarrival and document size distributions (rather than general distributions), which lead to some simplifications in our proofs. A particularly interesting aspect here is that in contrast to prior results on state space collapse for open multiclass queueing networks, our lifting map can be nonlinear (for $\alpha \neq 1$). In addition, unlike Assumption 3.1 in [5], we do not require an exponential rate of convergence of fluid model solutions towards points on the invariant manifold; we only use uniform convergence of fluid model solutions towards the invariant manifold (starting in a compact set). Despite these differences, our main line of argument follows that of Bramson [5].

We first provide some preliminary results in Section 6.1. Our proof of multiplicative state space collapse is then given in Section 6.2.

We note (for the extension to multi-path routing described in Section 5.5) that the proofs in this and the next section use that the entries in the matrix A are non-negative, rather than the stronger condition that they are zeros or ones.

6.1. Crucial Estimates for Fluid Scaled Processes. Recall the definition of the fluid scaled processes from Section 3. For each $r > 0$ and $m \in \mathbb{Z}_+$, let

$$(75) \quad \bar{n}_m^r = \lfloor \bar{N}^r(m) \rfloor \vee 1,$$

and define shifted and scaled processes $\bar{E}^{r,m}$, $\bar{S}^{r,m}$, $\bar{T}^{r,m}$, $\bar{U}^{r,m}$, $\bar{N}^{r,m}$, $\bar{W}^{r,m}$ as follows. For $i \in \mathbb{I}$, $j \in \mathbb{J}$, and $t \geq 0$,

$$(76) \quad \bar{E}_i^{r,m}(t) = \frac{\bar{E}_i^r(m + \bar{n}_m^r t) - \bar{E}_i^r(m)}{\bar{n}_m^r},$$

$$(77) \quad \bar{S}_i^{r,m}(t) = \frac{\bar{S}_i^r(\bar{T}_i^r(m) + \bar{n}_m^r t) - \bar{S}_i^r(\bar{T}_i^r(m))}{\bar{n}_m^r},$$

$$(78) \quad \bar{T}_i^{r,m}(t) = \frac{\bar{T}_i^r(m + \bar{n}_m^r t) - \bar{T}_i^r(m)}{\bar{n}_m^r},$$

$$(79) \quad \bar{U}_j^{r,m}(t) = \frac{\bar{U}_j^r(m + \bar{n}_m^r t) - \bar{U}_j^r(m)}{\bar{n}_m^r},$$

$$(80) \quad \bar{N}_i^{r,m}(t) = \frac{\bar{N}_i^r(m + \bar{n}_m^r t)}{\bar{n}_m^r},$$

$$(81) \quad \bar{W}_j^{r,m}(t) = \frac{\bar{W}_j^r(m + \bar{n}_m^r t)}{\bar{n}_m^r} = \sum_{i \in \mathbb{I}} A_{ji}(\mu_i^r)^{-1} \bar{N}_i^{r,m}(t).$$

The additional scaling by \bar{n}_m^r used here is to ensure that the starting values of $\bar{N}^{r,m}$ all lie in a single compact set, namely the unit ball in $\mathbb{R}_+^{\mathbb{I}}$. This facilitates use of the properties of fluid model solutions described in Section 4. It is easy to check using (5)–(6), (13)–(15), and the scaling property of Λ , that for each $i \in \mathbb{I}$, $j \in \mathbb{J}$, and $t \geq 0$,

$$(82) \quad \bar{N}_i^{r,m}(t) = \bar{N}_i^{r,m}(0) + \bar{E}_i^{r,m}(t) - \bar{S}_i^{r,m}(\bar{T}_i^{r,m}(t)),$$

$$(83) \quad \bar{U}_j^{r,m}(t) = C_j t - \sum_{i \in \mathbb{I}} A_{ji} \bar{T}_i^{r,m}(t),$$

$$(84) \quad \bar{T}_i^{r,m}(t) = \int_0^t \Lambda_i(\bar{N}^{r,m}(s)) ds.$$

The following theorem summarizes essential properties of the above processes needed for our proof of multiplicative state space collapse. The proof of this uses very similar arguments to those in Sections 4, 5 and 6 of Bramson [5].

THEOREM 6.1. *Let*

$$(85) \quad K = (1 + |\mu| \vee |\nu|)(1 + \max_{j \in \mathbb{J}} C_j).$$

Fix $T > 0$. For each $L \geq 1$ there is a sequence of measurable sets $\{\mathcal{G}_L^r : r > 0\}$ and a family of positive constants $\{r_{\varepsilon,L} : \varepsilon \in (0, 1)\}$ such that for each $\varepsilon \in (0, 1)$ and $r \geq r_{\varepsilon,L}$,

(i) $P(\mathcal{G}_L^r) \geq 1 - \varepsilon$,

(ii) for each $\omega \in \mathcal{G}_L^r$ and $m = 0, 1, \dots, \lfloor rT \rfloor$,

$$(86) \quad |\bar{N}^{r,m}(t, \omega) - \bar{N}^{r,m}(s, \omega)| \leq K|t - s| + \varepsilon \text{ for all } s, t \in [0, L],$$

and there is a fluid model solution $\tilde{n}(\cdot)$ satisfying

$$(87) \quad \|\bar{N}^{r,m}(\cdot, \omega) - \tilde{n}(\cdot)\|_L < \varepsilon.$$

Note that we have not indicated explicit dependence on T in the above as T will always be fixed in the application of this result. Also, to simplify notation, we have omitted explicit indication of the dependence of $\tilde{n}(\cdot)$ on r, m, ω, L and ε .

Before proving this theorem, we establish the following preliminary lemma. For this lemma and the proof of Theorem 6.1, let $\tilde{C} = \max_{j \in \mathbb{J}} C_j$.

LEMMA 6.1. *Fix $T > 0$. For each $L \geq 1$ there exists a sequence of measurable sets $\{\mathcal{G}_L^r : r > 0\}$ and a family of positive constants $\{r_{\varepsilon,L} : \varepsilon \in (0, 1)\}$ such that for each $\varepsilon \in (0, 1)$ and $r \geq r_{\varepsilon,L}$ we have*

$$(88) \quad P(\mathcal{G}_L^r) \geq 1 - \varepsilon,$$

and on \mathcal{G}_L^r , for $m = 0, 1, \dots, \lfloor rT \rfloor$,

$$(89) \quad \|\bar{E}^{r,m}(\cdot) - v^r(\cdot)\|_L \leq \varepsilon/4,$$

$$(90) \quad \|\bar{S}^{r,m}(\bar{T}^{r,m}(\cdot)) - \mu^r(\bar{T}^{r,m}(\cdot))\|_L \leq \varepsilon/4,$$

where $v^r(t) = v^r t$, $(\bar{S}^{r,m}(\bar{T}^{r,m}(t)))_i = \bar{S}_i^{r,m}(\bar{T}_i^{r,m}(t))$ and $(\mu^r(\bar{T}^{r,m}(t)))_i = \mu_i^r \bar{T}_i^{r,m}(t)$, for all $t \geq 0$ and $i \in \mathbb{I}$.

PROOF. Fix $L \geq 1$. Note that since the bandwidth allocations given by $\Lambda(\cdot)$ are bounded by \tilde{C} , for each $r > 0$ and $i \in \mathbb{I}$, $T_i^r(\cdot)$ is Lipschitz continuous with Lipschitz constant \tilde{C} . Since this property is unchanged by the (fluid) scaling in (78), we have that $\bar{T}_i^{r,m}(\cdot)$ is also Lipschitz continuous with the same Lipschitz constant. It follows from this that $\|\bar{T}^{r,m}(\cdot)\|_L \leq \mathbf{I}\tilde{C}L$. On combining this with Assumption 3.1, the fact that the interarrival times and document sizes are exponential, and the memoryless property of the associated Poisson processes, by a similar argument to that used in verifying (5.19) of Proposition 5.1 in Bramson [5], we have that as $r \rightarrow \infty$,

$$(91) \quad \max_{m=0}^{\lfloor rT \rfloor} (\|\bar{S}^{r,m}(\bar{T}^{r,m}(\cdot)) - \mu^r(\bar{T}^{r,m}(\cdot))\|_L \vee \|\bar{E}^{r,m}(\cdot) - v^r(\cdot)\|_L) \Rightarrow 0.$$

It follows that there exists a sequence $\{a_\ell\}_{\ell=0}^\infty$ satisfying $a_0 = 0$, $a_\ell \in (\ell \vee a_{\ell-1}, \infty)$ for each integer $\ell \geq 1$, such that for each $r \geq a_\ell$,

$$P \left[\bigcap_{m=0}^{\lfloor rT \rfloor} \left\{ \|\bar{S}^{r,m}(\bar{T}^{r,m}(\cdot)) - \mu^r(\bar{T}^{r,m}(\cdot))\|_L \vee \|\bar{E}^{r,m}(\cdot) - v^r(\cdot)\|_L \leq \frac{1}{2^{\ell+2}} \right\} \right] \geq 1 - \frac{1}{2^\ell}.$$

For $r > 0$, define

$$\delta(r) = \sum_{\ell=0}^{\infty} \frac{1}{2^\ell} \mathbf{1}_{[a_\ell, a_{\ell+1})}(r),$$

and

$$\mathcal{G}_L^r = \bigcap_{m=0}^{\lfloor rT \rfloor} \left\{ \|\bar{S}^{r,m}(\bar{T}^{r,m}(\cdot)) - \mu^r(\bar{T}^{r,m}(\cdot))\|_L \vee \|\bar{E}^{r,m}(\cdot) - v^r(\cdot)\|_L \leq \frac{\delta(r)}{4} \right\}.$$

Then, since $\delta(r) \rightarrow 0$ as $r \rightarrow \infty$, for each $\varepsilon \in (0, 1)$, there is $r_{\varepsilon,L} > 0$ such that $\delta(r) \leq \varepsilon$ for all $r \geq r_{\varepsilon,L}$. Combining the above we conclude that for all $r \geq r_{\varepsilon,L}$,

$$(92) \quad P(\mathcal{G}_L^r) \geq 1 - \delta(r) \geq 1 - \varepsilon. \quad \square$$

Proof of Theorem 6.1. Fix $T > 0$ and $L \geq 1$. Also, fix a sequence of measurable sets $\{\mathcal{G}_L^r : r > 0\}$ and a family of positive constants $\{r_{\varepsilon,L} : \varepsilon \in (0, 1)\}$ as in Lemma 6.1. Without loss of generality, by Assumption 3.1, we may further assume (by choosing $r_{\varepsilon,L}$ a bit larger if necessary) that for all $r \geq r_{\varepsilon,L}$,

$$|\nu^r| \leq 1 + |\nu| \quad \text{and} \quad |\mu^r| \leq 1 + |\mu|.$$

By Lemma 6.1, the uniform Lipschitz continuity of $\bar{T}^{r,m}(\cdot)$, and (82), for each $\varepsilon \in (0, 1)$ and $r \geq r_{\varepsilon,L}$ we have that on \mathcal{G}_L^r , for each $s, t \in [0, L]$ and $m = 0, 1, \dots, \lfloor rT \rfloor$,

$$(93) \quad |\bar{E}^{r,m}(t) - \bar{E}^{r,m}(s)| \leq \frac{\varepsilon}{2} + |\nu^r| |t - s|,$$

$$(94) \quad |\bar{S}^{r,m}(\bar{T}^{r,m}(t)) - \bar{S}^{r,m}(\bar{T}^{r,m}(s))| \leq \frac{\varepsilon}{2} + |\mu^r| \tilde{C} |t - s|,$$

$$(95) \quad |\bar{N}^{r,m}(t) - \bar{N}^{r,m}(s)| \leq \varepsilon + (|\nu^r| \vee |\mu^r|) (1 + \tilde{C}) |t - s|.$$

Then part (i) of Theorem 6.1 follows directly from Lemma 6.1. Inequality (86) follows from Lemma 6.1, (95) and the choice of $r_{\varepsilon,L}$ and K . As for the last part (87) of Theorem 6.1, for a proof by contradiction, suppose that there is an $\varepsilon_0 \in (0, 1)$ such that for each integer $\ell \geq r_{\varepsilon_0,L}$, there exists a value r_ℓ of r such that $r_\ell \geq \ell$, and $m_\ell \in \{0, 1, \dots, \lfloor r_\ell T \rfloor\}$, and $\omega_\ell \in \mathcal{G}_L^{r_\ell}$, such that for any fluid model solution \bar{n} we have

$$(96) \quad \|\bar{N}^{r_\ell, m_\ell}(\cdot, \omega_\ell) - \bar{n}(\cdot)\|_L \geq \varepsilon_0.$$

For the contradiction, we will show that there is a subsequence of $\{\bar{N}^{r_\ell, m_\ell}(\cdot, \omega_\ell)\}_{\ell=1}^\infty$ that converges uniformly on $[0, L]$ to a fluid model solution. For this, note that $\{\bar{N}^{r_\ell, m_\ell}(0, \omega_\ell)\}_{\ell=1}^\infty$ is bounded by one, and $\{\bar{T}^{r_\ell, m_\ell}(\cdot, \omega_\ell)\}_{\ell=1}^\infty$ is uniformly bounded and equicontinuous on $[0, L]$, by the uniform Lipschitz continuity of $\bar{T}^{r,m}$. Thus, using the Ascoli-Arzelà theorem, we have that along a subsequence, $\{\bar{N}^{r_\ell, m_\ell}(0, \omega_\ell)\}_{\ell=1}^\infty$ converges to a finite value and $\{\bar{T}^{r_\ell, m_\ell}(\cdot, \omega_\ell)\}_{\ell=1}^\infty$ converges uniformly on $[0, L]$. Fix such a convergent subsequence and denote the respective limits by $\bar{n}(0)$ and $\tau(\cdot)$. Then, using the convergence of $\{(\bar{E}^{r_\ell, m_\ell}(\cdot, \omega_\ell), \bar{S}^{r_\ell, m_\ell}(\bar{T}^{r_\ell, m_\ell}(\cdot, \omega_\ell)))\}_{\ell=1}^\infty$ implied by (89) and (90), we have that along the same subsequence, $\{\bar{N}^{r_\ell, m_\ell}(\cdot, \omega_\ell)\}_{\ell=1}^\infty$ converges uniformly on $[0, L]$ to $\bar{n}(\cdot)$ defined by

$$(97) \quad \bar{n}_i(t) = \bar{n}_i(0) + \nu_i t - \mu_i \tau_i(t), \quad t \in [0, L], \quad i \in \mathbb{I}.$$

Furthermore, by following the pathwise argument in Appendix B of Kelly and Williams [26], one can show that at each $t \in (0, L]$ where the derivative of $\bar{n}(\cdot)$ exists (this means the left hand derivative at $t = L$), we have that (22) and (23) hold with \bar{n} in place of n there. It follows that $\bar{n}(\cdot)$ has the same properties as a fluid model solution on $[0, L]$. We can easily extend \bar{n} to be a fluid model solution

on $[0, \infty)$ by defining it on (L, ∞) to equal $\bar{f}(\cdot - L)$, where $\bar{f} : [0, \infty) \rightarrow \mathbb{R}_+^{\mathbf{I}}$ is a fluid model solution satisfying $\bar{f}(0) = \bar{n}(L)$; the existence of such a fluid model solution follows from Theorem B.1 in [26]. This yields the desired contradiction of (96). \square

6.2. Proof of Multiplicative State Space Collapse. We will use Theorem 6.1 to show that for $T > 0$ fixed, for each $\bar{L} \geq 1$, there is $L > \bar{L}$ such that for each $\varepsilon \in (0, 1)$, for all r sufficiently large, with probability at least $1 - \varepsilon$ the left member of (33) is dominated by

$$(98) \quad \left\| \bar{N}^{r,0}(\cdot) - \Delta(\bar{W}^{r,0}(\cdot)) \right\|_{[0, \bar{L}]} + \sup_{m=0}^{\lfloor rT \rfloor - 1} \left\| \bar{N}^{r,m}(\cdot) - \Delta(\bar{W}^{r,m}(\cdot)) \right\|_{[\bar{L}, L]}.$$

where the quantities $\{\bar{N}^{r,m}, m = 0, 1, \dots, \lfloor rT \rfloor - 1\}$ can be approximated by fluid model solutions over $[0, L]$. The results of Section 4 on the behavior of fluid model solutions, especially Proposition 4.4, can then be combined with a suitable choice of \bar{L} and the assumptions of Theorem 5.1 on the initial conditions, to prove multiplicative state space collapse. We now give the detailed proof.

Proof of Theorem 5.1. Fix $T > 0$. For each $r > 0$, let \mathcal{H}_T^r denote the left member of (33). Since for each $r > 0$ and $t \geq 0$, we have

$$(99) \quad \hat{N}^r(t) = \bar{N}^r(rt), \quad \hat{W}^r(t) = \bar{W}^r(rt) = A(M^r)^{-1} \bar{N}^r(rt),$$

it follows that

$$(100) \quad \mathcal{H}_T^r = \frac{\|\bar{N}^r(\cdot) - \Delta(\bar{W}^r(\cdot))\|_{rT}}{\|\bar{N}^r(\cdot)\|_{rT} \vee 1}.$$

For $\bar{L} \geq 1$ fixed, intervals of the form $[m, m + \bar{L}\bar{n}_m^r], m = 0, 1, \dots, \lfloor rT \rfloor$ cover $[0, rT]$ since $\bar{n}_m^r \geq 1$ for each $m \in \mathbb{Z}_+$. Hence, for each $r > 0$ and $t \in [0, rT]$, the following is well defined as a random variable taking values in $\{0, 1, \dots, \lfloor rT \rfloor\}$:

$$(101) \quad m^r(t) = \inf\{m : m \leq t \leq m + \bar{L}\bar{n}_m^r\}.$$

For any $r > 0$ and $t \in [0, rT]$ such that $t > \bar{L}\bar{n}_0^r$, we have $m^r(t) \geq 1$ and then

$$(102) \quad t - (m^r(t) - 1) > \bar{L}\bar{n}_{m^r(t)-1}^r,$$

for otherwise $m^r(t)$ could be replaced by a value of m strictly less than $m^r(t)$. Thus, for each $r > 0$ and $t \in [0, rT]$ such that $t > \bar{L}\bar{n}_0^r$, we have

$$(103) \quad m^r(t) - 1 + \bar{L}\bar{n}_{m^r(t)-1}^r < t \leq m^r(t) + \bar{L}\bar{n}_{m^r(t)}^r,$$

where $m^r(t) \in \{1, 2, \dots, \lfloor rT \rfloor\}$. From (100) and (103) we have that \mathcal{H}_T^r is dominated by

$$(104) \quad \frac{\|\overline{N}^r(\cdot) - \Delta(\overline{W}^r(\cdot))\|_{[0, \overline{L}\overline{n}_0^r]}}{\overline{n}_0^r} + \sup_{m=1}^{\lfloor rT \rfloor} \frac{\|\overline{N}^r(\cdot) - \Delta(\overline{W}^r(\cdot))\|_{[m-1+\overline{L}\overline{n}_{m-1}^r, m+\overline{L}\overline{n}_m^r]}}{\overline{n}_{m-1}^r} \\ = \left\| \overline{N}^{r,0}(\cdot) - \Delta(\overline{W}^{r,0}(\cdot)) \right\|_{[0, \overline{L}]} + \sup_{m=1}^{\lfloor rT \rfloor} \left\| \overline{N}^{r,m-1}(\cdot) - \Delta(\overline{W}^{r,m-1}(\cdot)) \right\|_{\left[\overline{L}, \frac{1+\overline{L}\overline{n}_m^r}{\overline{n}_{m-1}^r}\right]},$$

where we have used the scaling property of $\Delta(\cdot)$ in obtaining the equality (cf. Proposition 4.1).

Recall the definition of the constant $K \geq 1$ from (85). Given $\overline{L} \geq 1$, let $L = 4K\overline{L}$. Below we refer to the sequence of measurable sets $\{\mathcal{G}_L^r : r > 0\}$ introduced in Theorem 6.1. Focusing on the right endpoints of the time intervals appearing in the last term of (104), we next show that, for all r sufficiently large, on \mathcal{G}_L^r , for $m = 1, \dots, \lfloor rT \rfloor$, we have

$$(105) \quad \frac{1 + \overline{L}\overline{n}_m^r}{\overline{n}_{m-1}^r} \leq L.$$

For this, note that by (80), we have for $m = 1, 2, \dots$,

$$(106) \quad \frac{\overline{N}^r(m)}{\overline{n}_{m-1}^r} = \frac{\overline{N}^r(m-1+1)}{\overline{n}_{m-1}^r} = \overline{N}^{r,m-1}\left(\frac{1}{\overline{n}_{m-1}^r}\right).$$

On the other hand, by (75), for $m = 1, 2, \dots$,

$$(107) \quad \left| \overline{N}^{r,m-1}(0) \right| = \left| \frac{\overline{N}^r(m-1)}{\overline{n}_{m-1}^r} \right| \leq 1.$$

It follows from (106), (107) and Theorem 6.1 that for each $\varepsilon \in (0, 1)$, for each $r \geq r_{\varepsilon, L}$ and $m = 1, \dots, \lfloor rT \rfloor$, on \mathcal{G}_L^r we have

$$(108) \quad \frac{\overline{n}_m^r}{\overline{n}_{m-1}^r} = \frac{|\overline{N}^r(m)|}{\overline{n}_{m-1}^r} \vee \frac{1}{\overline{n}_{m-1}^r} \\ \leq \left| \overline{N}^{r,m-1}\left(\frac{1}{\overline{n}_{m-1}^r}\right) \right| \vee 1 \\ \leq \left(\left| \overline{N}^{r,m-1}(0) \right| + K \left(\frac{1}{\overline{n}_{m-1}^r} \right) + \varepsilon \right) \vee 1 \\ \leq (1 + K \cdot 1 + 1) \vee 1 \\ \leq 3K,$$

and hence (105) holds. On combining (104) with (105), we have that for each $\varepsilon \in (0, 1)$, for all $r \geq r_{\varepsilon, L}$, the following inequality holds on \mathcal{G}_L^r ,

$$(109) \quad \begin{aligned} \mathcal{H}_T^r &\leq \left\| \overline{N}^{r,0}(\cdot) - \Delta(\overline{W}^{r,0}(\cdot)) \right\|_{[0, L]} \\ &\quad + \sup_{m=0}^{\lfloor rT \rfloor - 1} \left\| \overline{N}^{r,m}(\cdot) - \Delta(\overline{W}^{r,m}(\cdot)) \right\|_{[\overline{L}, L]}. \end{aligned}$$

We will now estimate the two terms on the right hand side of (109). In brief, the idea is to use Theorem 6.1 to show that for all r sufficiently large, for each $\omega \in \mathcal{G}_L^r$, $m = 0, 1, \dots, \lfloor rT \rfloor$, there is a fluid model solution \tilde{n} that is uniformly close to $\overline{N}^{r,m}(\cdot, \omega)$ over the time interval $[0, L]$. For the first term in (109), we then use the assumptions on the initial conditions and Proposition 4.5 to show that for all r sufficiently large, with high probability, \tilde{n} is uniformly close to the invariant manifold \mathcal{M}_α on $[0, \infty)$ and use this to control the first term in (109). For the last term in (109), we use Proposition 4.4 to choose \overline{L} (independent of r, ω or m), so that \tilde{n} is uniformly close to the invariant manifold \mathcal{M}_α on $[\overline{L}, \infty)$, and so for all r sufficiently large, $\omega \in \mathcal{G}_L^r$, and $m = 0, 1, \dots, \lfloor rT \rfloor - 1$, for each $t \in [\overline{L}, L]$, there is a point \tilde{n}_t^* (depending on r, ω, m, t) in \mathcal{M}_α that is uniformly close to $\overline{N}^{r,m}(t, \omega)$. Then,

$$\begin{aligned} &|\overline{N}^{r,m}(t, \omega) - \Delta(\overline{W}^{r,m}(t, \omega))| \\ &\leq |\overline{N}^{r,m}(t, \omega) - \tilde{n}_t^*| + |\tilde{n}_t^* - \Delta(AM^{-1}\tilde{n}_t^*)| \\ &\quad + |\Delta(AM^{-1}\tilde{n}_t^*) - \Delta(A(M^r)^{-1}\overline{N}^{r,m}(t, \omega))|, \end{aligned}$$

where each of the last three terms can be made small (uniformly for $\omega \in \mathcal{G}_L^r, m = 0, 1, \dots, \lfloor rT \rfloor - 1, t \in [\overline{L}, L]$, for all r sufficiently large), by the choice of \tilde{n}_t^* , the fact that $n = \Delta(w(n))$ for a point n on the invariant manifold \mathcal{M}_α , the continuity of Δ , and the convergence of $(M^r)^{-1}$ to M^{-1} as $r \rightarrow \infty$. We now give the full details of the argument. In the first paragraph below we develop estimates that will be used for both the first and the second term on the right hand side of (109). (Accordingly, we consider values of m that include $\lfloor rT \rfloor$, as this value may be zero.)

Fix $\eta \in (0, 1)$. (In this proof only, we locally reuse the symbol η for a positive constant in $(0, 1)$. This is distinct from the use of η elsewhere as a probability measure.) By Proposition 4.3, there is a constant $D(9/8) \geq 9/8$ such that any fluid model solution $n(\cdot)$ satisfying $|n(0)| \leq 9/8$ satisfies $|n(t)| \leq D(9/8)$ for all $t \geq 0$. By the uniform continuity of $\Delta(\cdot)$ on compact sets in \mathbb{R}_+^J (cf. Proposition 4.1) and the fact that $(M^r)^{-1} \rightarrow M^{-1}$ as $r \rightarrow \infty$ (cf. Assumption 3.1), there are constants $r_\eta > 0$ and $\gamma \in (0, \eta/4)$ such that for all $r \geq r_\eta$,

$$(110) \quad \left| \Delta(AM^{-1}n) - \Delta(A(M^r)^{-1}n') \right| < \frac{\eta}{8}$$

whenever $n, n' \in \mathbb{R}_+^I$, $|n| \vee |n'| \leq 1 + D(9/8)$ and $|n - n'| < 2\gamma$. By Proposition 4.5, there is $\varepsilon \in (0, \gamma/2)$ such that if $n(\cdot)$ is a fluid model solution satisfying $|n(0)| \leq 9/8$ and $d(n(0), \mathcal{M}_\alpha) < 2\varepsilon$, then

$$(111) \quad d(n(t), \mathcal{M}_\alpha) < \gamma \quad \text{for all } t \geq 0.$$

Let $\bar{L} = T_{\frac{9}{8}, \varepsilon}$, where $T_{R, \varepsilon}$ is defined in Proposition 4.4. Set $L = 4K\bar{L}$ as above. Then, by Theorem 6.1, for each $r \geq r_{\varepsilon, L}$, $m = 0, 1, \dots, \lfloor rT \rfloor$, and $\omega \in \mathcal{G}_L^r$, there is a fluid model solution $\tilde{n}(\cdot)$ such that

$$(112) \quad \|\bar{N}^{r, m}(\cdot, \omega) - \tilde{n}(\cdot)\|_L < \varepsilon.$$

We fix such r, m , and ω for the remainder of this paragraph. Since $|\bar{N}^{r, m}(0, \omega)| \leq 1$ and $\varepsilon < 1/8$, we have $|\tilde{n}(0)| < 9/8$. It then follows from Proposition 4.4 and the choice of \bar{L} that

$$(113) \quad d(\tilde{n}(t), \mathcal{M}_\alpha) < \varepsilon \quad \text{for all } t \geq \bar{L}.$$

For each $t \in [\bar{L}, L]$, let $\tilde{n}_t^* \in \mathcal{M}_\alpha$ such that

$$(114) \quad |\tilde{n}(t) - \tilde{n}_t^*| < \varepsilon.$$

On combining this with (112), we obtain for each $t \in [\bar{L}, L]$,

$$(115) \quad |\bar{N}^{r, m}(t, \omega) - \tilde{n}_t^*| < 2\varepsilon < \gamma.$$

By Proposition 4.3 and the fact that $|\tilde{n}(0)| \leq 9/8$, we have $|\tilde{n}(t)| \leq D(9/8)$ for all $t \geq 0$. So, by (114) and (112), for each $t \in [\bar{L}, L]$,

$$|\tilde{n}_t^*| \leq \varepsilon + D(9/8) < 1 + D(9/8)$$

and

$$|\bar{N}^{r, m}(t, \omega)| \leq \varepsilon + D(9/8) < 1 + D(9/8).$$

On combining the above and using the fact that $\bar{W}^{r, m} = A(M^r)^{-1}\bar{N}^{r, m}$, we have for each $r \geq r_{\varepsilon, L} \vee r_\eta$, $m = 0, 1, \dots, \lfloor rT \rfloor - 1$, $\omega \in \mathcal{G}_L^r$ and $t \in [\bar{L}, L]$,

$$(116) \quad \begin{aligned} & |\bar{N}^{r, m}(t, \omega) - \Delta(\bar{W}^{r, m}(t, \omega))| \\ & \leq |\bar{N}^{r, m}(t, \omega) - \tilde{n}_t^*| + |\tilde{n}_t^* - \Delta(AM^{-1}\tilde{n}_t^*)| \\ & \quad + |\Delta(AM^{-1}\tilde{n}_t^*) - \Delta(A(M^r)^{-1}\bar{N}^{r, m}(t, \omega))| \\ & \leq 2\varepsilon + 0 + \frac{\eta}{8} < \frac{\eta}{2}. \end{aligned}$$

For the second inequality, we have used (115), part (iv) of Theorem 4.1, and (110). This takes care of estimating the last term in (109).

To estimate the first term on the right hand side of (109), we need to use the initial behavior of \bar{N}^r . By the assumption in the theorem,

$$(117) \quad \left| \bar{N}^{r,0}(0) - \Delta(\bar{W}^{r,0}(0)) \right| = |\hat{N}^r(0) - \Delta(\hat{W}^r(0))| \rightarrow 0$$

in probability as $r \rightarrow \infty$. By Proposition 4.2 we have $\Delta(\bar{W}^{r,0}(0)) \in \mathcal{M}_\alpha$, and so it follows that

$$(118) \quad d(\bar{N}^{r,0}(0), \mathcal{M}_\alpha) \rightarrow 0 \quad \text{in probability as } r \rightarrow \infty.$$

Let $r_\varepsilon > 0$ such that

$$(119) \quad P\left(d(\bar{N}^{r,0}(0), \mathcal{M}_\alpha) \geq \varepsilon\right) < \varepsilon \quad \text{for all } r \geq r_\varepsilon.$$

For each $r \geq r_\varepsilon \vee r_{\varepsilon,L}$ and $\omega \in \mathcal{G}_L^r$ satisfying $d(\bar{N}^{r,0}(0, \omega), \mathcal{M}_\alpha) < \varepsilon$, by (112) we have

$$(120) \quad |\tilde{n}(0)| \leq 1 + \varepsilon \leq \frac{9}{8} \quad \text{and} \quad d(\tilde{n}(0), \mathcal{M}_\alpha) < \varepsilon + \varepsilon = 2\varepsilon.$$

For such r, ω , it follows from (111), that for each $t \in [0, \infty)$ there is $\tilde{n}_t^* \in \mathcal{M}_\alpha$ such that $|\tilde{n}(t) - \tilde{n}_t^*| < \gamma$. Hence by (112), for $t \in [0, \bar{L}] \subset [0, L]$, we have

$$(121) \quad |\bar{N}^{r,0}(t, \omega) - n_t^*| < \varepsilon + \gamma < \frac{3\gamma}{2}.$$

Then, in a similar manner to that used in showing (116), we have for all $t \in [0, \bar{L}]$, $r \geq r_\varepsilon \vee r_{\varepsilon,L} \vee r_\eta$, and $\omega \in \mathcal{G}_L^r$ satisfying $d(\bar{N}^{r,0}(0, \omega), \mathcal{M}_\alpha) < \varepsilon$,

$$(122) \quad |\bar{N}^{r,0}(t, \omega) - \Delta(\bar{W}^{r,0}(t, \omega))| \leq (\varepsilon + \gamma) + 0 + \frac{\eta}{8} < \frac{\eta}{2}.$$

By combining Theorem 6.1, (109), (116), (119) and (122), we have that for all $r \geq r_\varepsilon \vee r_{\varepsilon,L} \vee r_\eta$,

$$(123) \quad \begin{aligned} P(\mathcal{H}_T^r \geq \eta) &\leq P\left(\{\mathcal{H}_T^r \geq \eta\} \cap \mathcal{G}_L^r \cap \left\{d(\bar{N}^{r,0}(0), \mathcal{M}_\alpha) < \varepsilon\right\}\right) \\ &\quad + P((\mathcal{G}_L^r)^c) + P\left(d(\bar{N}^{r,0}(0), \mathcal{M}_\alpha) \geq \varepsilon\right) \\ &\leq 0 + \varepsilon + \varepsilon = 2\varepsilon < \frac{\eta}{4}. \end{aligned}$$

Since $\eta \in (0, 1)$ was arbitrary, it follows that $\mathcal{H}_T^r \rightarrow 0$ in probability as $r \rightarrow \infty$. \square

7. PROOF OF DIFFUSION APPROXIMATION WHEN $\alpha = 1$. Throughout this section we assume that $\alpha = 1$ and that the local traffic condition, Assumption 5.1, holds. To prove the diffusion approximation result, Theorem 5.2, we shall use the invariance principle in [20]. A key assumption for that theory to yield convergence (rather than just C -tightness) is that there is existence and uniqueness in law for the limit diffusion process \tilde{W} , which follows in the case $\alpha = 1$ from work of Dai and Williams [10]. We first verify that the basic assumptions of [20] and [10] are satisfied by the state space \mathscr{W}_1 (cf. (47)) and directions of reflection $\{\gamma^j : j \in \mathbb{J}\}$. For this we need the following definition. Also, recall the definition of the matrix B following (47).

DEFINITION 7.1. *A $d \times d$ matrix D is completely- S if and only if for each principal submatrix \tilde{D} of D there is a vector $\tilde{x} \geq 0$ such that $\tilde{D}\tilde{x} > 0$. (Here a principal submatrix of D is a matrix obtained by deleting all rows and columns of D with indices in some strict subset of $\{1, 2, \dots, d\}$.)*

LEMMA 7.1. *The $\mathbf{J} \times \mathbf{J}$ symmetric matrix ABA' is positive definite and invertible. The inverse matrix $(ABA')^{-1}$ is positive definite and completely- S .*

PROOF. Since A has full row rank and B is strictly positive definite, then the symmetric matrix ABA' is also strictly positive definite. Indeed, for a nonzero vector $x \in \mathbb{R}^{\mathbf{J}}$, $x'ABA'x = (A'x)'B(A'x)$. Since A has full row rank, $A'x$ is also nonzero and then since B is strictly positive definite we have that $(A'x)'B(A'x) > 0$. Since ABA' is strictly positive definite, it is invertible and the inverse matrix $(ABA')^{-1}$ is also strictly positive definite. It then follows that $(ABA')^{-1}$ is a P -matrix, i.e., all principal minors are positive, and hence $(ABA')^{-1}$ is completely- S (see [9], especially Theorems 3.3.7, 3.9.11 and Corollary 3.9.13 for the relationship between P -matrices and completely- S matrices). \square

The results we need to apply from [20] to prove Theorem 5.2 require that Assumptions **(A1)–(A5)** and 5.1 of [20] hold. We shall not fully describe these general conditions. However, we shall now indicate the meaning of these assumptions for the context treated here. Assumptions **(A1)–(A3)** are restrictions on the state space and Assumptions **(A4)–(A5)** pertain to the directions of reflection, all for the limit diffusion. As noted in [20], Assumptions **(A1)–(A3)** are satisfied by a convex polyhedron with non-empty interior that is described as the intersection of a minimal set of half-spaces. As we shall see in the next lemma, the state space \mathscr{W}_1 has this convex polyhedral form. Assumption **(A4)** requires that the reflection vector field on each boundary face be uniformly Lipschitz continuous and of unit length. Since our reflection vectors are constant unit vectors on each boundary face, this condition is trivially satisfied in our context. Assumption **(A5)** imposes geometric

conditions on the directions of reflection. These conditions are generalizations for domains with piecewise smooth boundaries of conditions identified earlier by Dai and Williams [10] for existence and uniqueness of SRBMs living in convex polyhedral domains with a constant direction of reflection on each boundary face. The conditions of [10] are labelled as Assumption 1.1 in [10] and as Assumption 5.1 in [20]. (We shall use the latter label here.) For simple convex polyhedral cones (as we have here with \mathcal{W}_1), this assumption can be expressed as a completely- \mathcal{S} condition on a suitable matrix formed using the directions of reflection and the normals to the boundary faces. We will show that this condition is satisfied in our context, and as a consequence Assumption (A5) of [20] will follow immediately. We now formally verify that the aforementioned assumptions all hold in our context.

LEMMA 7.2. *Assumptions (A1)–(A5) and Assumption 5.1 of [20] are all satisfied by \mathcal{W}_1 and $\{\gamma^j : j \in \mathbb{J}\}$.*

PROOF. For $j \in \mathbb{J}$, let n^j be the vector given by the j^{th} row of $(ABA')^{-1}$. Since $(ABA')^{-1}$ is symmetric, n^j is also the j^{th} column of $(ABA')^{-1}$. By (47) and (45), we have

$$(124) \quad \mathcal{W}_1 = \{w \in \mathbb{R}^{\mathbb{J}} : n^j \cdot w \geq 0 \text{ for all } j \in \mathbb{J}\},$$

and

$$(125) \quad \mathcal{W}_1^j = \{ABA'q : q \in \mathbb{R}_+^{\mathbb{J}}, q_j = 0\} = \{w \in \mathcal{W}_1 : n^j \cdot w = 0\}, \quad j \in \mathbb{J}.$$

Since the $n^j, j \in \mathbb{J}$, are linearly independent, it follows that \mathcal{W}_1 is a simple convex polyhedral cone with minimal representation given by the intersection of the half-spaces $\{w \in \mathbb{R}^{\mathbb{J}} : n^j \cdot w \geq 0\}, j \in \mathbb{J}$. It is easy to see that \mathcal{W}_1 has non-empty interior. As noted in [20], it follows that Assumptions (A1)–(A3) of [20] are satisfied. Since the $\{\gamma^j, j \in \mathbb{J}\}$ are constant vectors, they are trivially uniformly Lipschitz continuous and so Assumption (A4) of [20] holds. As noted in Section 5.2 of [20], Assumption (A5) will hold if Assumption 5.1 of [20] holds. For this, since \mathcal{W}_1 is a simple convex polyhedron, it is equivalent to verify that the matrix NR is completely- \mathcal{S} , where N is the $\mathbb{J} \times \mathbb{J}$ matrix whose rows are given by the normals $n^j, j \in \mathbb{J}$, and R is the matrix whose columns are given by the vectors $\gamma^j, j \in \mathbb{J}$. (In fact, N is normalized so that the rows have unit length in [20], but one may equivalently use the unnormalized matrix N .) Now, $N = (ABA')^{-1}$ and $R = I$, the $\mathbb{J} \times \mathbb{J}$ identity matrix, and so $NR = (ABA')^{-1}$ which we know is completely- \mathcal{S} by Lemma 7.1. \square

Throughout this section, we shall need various constants in inequalities. Some of these constants are denoted by \hat{C}_j , for $j = 0, 1, 2, \dots$. The hat has been added to the

notation here simply to distinguish these constants from the bandwidth capacities C_j . The next two lemmas provide basic ingredients for the proof of Theorem 5.2.

LEMMA 7.3. *There exist constants $\hat{C}_1 > 0$, $\hat{C}_2 \geq 1$ and $\bar{r} > 0$ such that for each $T > 0$,*

$$(126) \quad \hat{C}_1 \|\hat{W}^r(\cdot)\|_T \leq \|\hat{N}^r(\cdot)\|_T \leq \hat{C}_2 \|\hat{W}^r(\cdot)\|_T \quad \text{for all } r > \bar{r}.$$

PROOF. By Assumption 3.1, there is $\bar{r} > 0$ such that $\|A(M^r)^{-1}\| \leq 2\|AM^{-1}\|$ and $|\mu^r| < 2|\mu|$ for all $r > \bar{r}$. For each $t \geq 0$, by (17), we have that

$$(127) \quad |\hat{W}^r(t)| \leq \|A(M^r)^{-1}\| |\hat{N}^r(t)|.$$

Now, $\|AM^{-1}\| > 0$ and so by letting $\hat{C}_1 = (2\|AM^{-1}\|)^{-1}$ we see that the left hand inequality in (126) holds for any $T > 0$, for all $r > \bar{r}$. On the other hand, for each $i \in \mathbb{I}$, there is at least one $j_i \in \mathbb{J}$ such that $A_{j_i i} > 0$, and then using (17) again we obtain that for each $t \geq 0$,

$$A_{j_i i} (\mu_i^r)^{-1} \hat{N}_i^r(t) \leq \hat{W}_{j_i}^r(t).$$

By letting $\hat{C}_2 = \max(1, 2|\mu| \mathbf{I} \max_{i \in \mathbb{I}} (A_{j_i i})^{-1})$, we have that the right hand inequality in (126) holds for any $T > 0$ and $r > \bar{r}$. \square

LEMMA 7.4. *The sequence of processes $\{\hat{W}^r(0) + \hat{X}^r(\cdot), r > 0\}$ is C -tight.*

PROOF. Recall that for each $t \geq 0$,

$$(128) \quad \hat{X}^r(t) = A(M^r)^{-1} \left(\hat{E}^r(t) - \hat{S}^r(\bar{T}^r(t)) \right) + r(A\rho^r - C)t.$$

Since each route must use at least one resource, for each $n \in \mathbb{R}_+^{\mathbf{I}}$ and $i \in \mathbb{I}$, the bandwidth allocation $\Lambda_i(n)$ must be bounded by $\tilde{C} = \max_{j \in \mathbb{J}} C_j$. It follows that for each $r > 0$, $\bar{T}^r(\cdot)$ is uniformly Lipschitz continuous with Lipschitz constant \tilde{C} , and hence $\{\bar{T}^r(\cdot)\}$ is C -tight. On combining this with the functional central limit assumption (21), the convergence of $\{r(A\rho^r - C)t, r > 0\}$ and $\{M^r, r > 0\}$, and the convergence assumption on $\hat{W}^r(0)$ made at the end of Section 3, the desired result follows. \square

With the above preliminaries in place, we are now ready to address the main part of the proof of Theorem 5.2. In the following, we shall reuse some notation for local proof purposes. In particular, ξ and ζ are used for different purposes here than earlier in the paper.

Recall from (36) that for each $r > 0$ and $t \geq 0$,

$$(129) \quad \hat{W}^r(t) = \hat{W}^r(0) + \hat{X}^r(t) + \hat{U}^r(t),$$

and that by (17),

$$(130) \quad \hat{W}^r(t) = A(M^r)^{-1} \hat{N}^r(t) = \tilde{W}^r(t) + \hat{\xi}^r(t),$$

where

$$(131) \quad \tilde{W}^r(t) = AM^{-1} \Delta(\hat{W}^r(t)),$$

$$(132) \quad \hat{\xi}^r(t) = AM^{-1}(\hat{N}^r(t) - \Delta(\hat{W}^r(t))) + A((M^r)^{-1} - M^{-1})\hat{N}^r(t).$$

By Proposition 4.2 and the definition of \mathscr{W}_1 from \mathscr{M}_1 , we have that

$$(133) \quad \tilde{W}^r(t) \in \mathscr{W}_1 \quad \text{for all } t \geq 0 \text{ and for all } r > 0.$$

Returning to (129), for $\delta > 0$ fixed and each $r > 0, t \geq 0, j \in \mathbb{J}$,

$$(134) \quad \hat{U}_j^r(t) = \int_0^t 1_{\{d(\tilde{W}^r(s), \mathscr{W}_1^j) \leq \delta\}} d\hat{U}_j^r(s) + \hat{\xi}_j^{r,\delta}(t),$$

where

$$(135) \quad \hat{\xi}_j^{r,\delta}(t) = \int_0^t 1_{\{d(\tilde{W}^r(s), \mathscr{W}_1^j) > \delta\}} d\hat{U}_j^r(s).$$

A main step in establishing the diffusion approximation result, Theorem 5.2, is to show that for each $j \in \mathbb{J}$, with probability tending to one as $r \rightarrow \infty$, the process \hat{U}^r increases only when \tilde{W}^r (or \hat{W}^r) is near the boundary portion \mathscr{W}_1^j . The local traffic condition (Assumption 5.1) is used in verifying this in the following lemma, which also shows that state space collapse holds and that for fixed $T > 0$, with high probability we can obtain a uniform bound on $\|\hat{W}^r(\cdot)\|_T$ and on $\|\hat{U}^r(\cdot)\|_T$.

LEMMA 7.5. *Suppose, in addition to the assumptions that $\alpha = 1$ and that Assumption 5.1 holds, we have that*

$$(136) \quad |\hat{N}^r(0) - \Delta(\hat{W}^r(0))| \rightarrow 0$$

in probability as $r \rightarrow \infty$. Then for each $T > 0, \delta > 0$, there exist constants $K(T, \delta) > 0$ and $r(T, \delta) > 0$ such that for each $r \geq r(T, \delta)$,

$$(137) \quad \begin{aligned} &P\left(\|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T \leq \delta, \|\hat{\xi}^r(\cdot)\|_T \leq \delta, \|\hat{\xi}_j^{r,\delta}(\cdot)\|_T = 0, \right. \\ &\left. \|\hat{W}^r(\cdot)\|_T \leq K(T, \delta), \|\hat{U}^r(\cdot)\|_T \leq K(T, \delta)\right) \geq 1 - \delta. \end{aligned}$$

A main aspect of the proof of Lemma 7.5 is to show that for fixed $T > 0$, with high probability, we can obtain a suitable uniform bound on $\|\hat{N}^r(\cdot)\|_T$ (or equivalently, on $\|\hat{W}^r(\cdot)\|_T$) for all r sufficiently large. We shall use an oscillation inequality for the proof of this. A local version of this inequality was established in Theorem 4.1 of Kang and Williams [20]. For $\alpha \neq 1$, one would need to use that local version. However, for the case $\alpha = 1$ treated here, one can choose ρ in Theorem 4.1 of [20] to be arbitrarily large and consequently obtain a global version of the oscillation inequality. Here, for the case of $\alpha = 1$ only, rather than using Theorem 4.1 of [20], we shall give an alternative proof of a global oscillation inequality by first invoking a linear transformation to transform \mathscr{W}_1 to the orthant and then applying an oscillation inequality developed earlier by Williams [43] for that state space. This inequality has a slightly simpler form than that which would follow from [20].

For the statement of the oscillation inequality, we need the following notation. For any $0 \leq s < t < \infty$, let $\mathbb{D}([s, t], \mathbb{R}^{\mathbf{J}})$ denote the set of functions $x : [s, t] \rightarrow \mathbb{R}^{\mathbf{J}}$ that are right continuous on $[s, t)$ and have finite left limits on $(s, t]$, and for $x \in \mathbb{D}([s, t], \mathbb{R}^{\mathbf{J}})$, let

$$(138) \quad \text{Osc}(x, [s, t]) = \sup\{|x(v) - x(u)| : s \leq u < v \leq t\}.$$

PROPOSITION 7.1. (*Oscillation Inequality*) *There exists a constant $\hat{C}_0 > 0$ such that for any $\delta \geq 0$ and any $0 \leq s < t < \infty$, $w, x, y \in \mathbb{D}([s, t], \mathbb{R}^{\mathbf{J}})$ satisfying*

- (i) $w(u) = x(u) + \sum_{j \in \mathbb{J}} y_j(u) \gamma^j$ for all $u \in [s, t]$,
- (ii) $w(u) \in \mathscr{W}_1$ for all $u \in [s, t]$,
- (iii) for each $j \in \mathbb{J}$,
 - (a) $y_j(s) \geq 0$,
 - (b) y_j is non-decreasing,
 - (c) $y_j(u) = y_j(s) + \int_{(s, u]} \mathbf{1}_{\{d(w(v), \mathscr{W}_1^j) \leq \delta\}} dy_j(v)$ for all $u \in [s, t]$,

then the following hold:

$$(139) \quad \text{Osc}(w, [s, t]) \leq \hat{C}_0(\text{Osc}(x, [s, t]) + \delta),$$

$$(140) \quad \text{Osc}(y, [s, t]) \leq \hat{C}_0(\text{Osc}(x, [s, t]) + \delta).$$

PROOF. For w, x, y satisfying (i)–(iii) above, let $\tilde{w} = (ABA')^{-1}w$, $\tilde{x} = (ABA')^{-1}x$, $\tilde{y} = y$, and let $R = (ABA')^{-1}$. Since ABA' is a bijection from $\mathbb{R}_+^{\mathbf{J}}$ onto \mathscr{W}_1 , the matrix with columns given by the vectors $\{\gamma^j, j \in \mathbb{J}\}$ is the $\mathbf{J} \times \mathbf{J}$ identity matrix, and for $z \in \mathscr{W}_1$ the distance $d(z, \mathscr{W}_1^j) = n^j \cdot z / |n^j|$ for each $j \in \mathbb{J}$ (where n^j is the vector given by the j^{th} row of $(ABA')^{-1}$), it follows that $\tilde{w}, \tilde{x}, \tilde{y}$ satisfy

- (i) $\tilde{w}(u) = \tilde{x}(u) + R\tilde{y}(u)$ for all $u \in [s, t]$,

- (ii) $\tilde{w}(u) \in \mathbb{R}_+^{\mathbf{J}}$ for all $u \in [s, t]$,
- (iii) for each $j \in \mathbb{J}$,
 - (a) $\tilde{y}_j(s) \geq 0$,
 - (b) \tilde{y}_j is non-decreasing,
 - (c) $\tilde{y}_j(u) = \tilde{y}_j(s) + \int_{(s,u]} 1_{\{\tilde{w}_j(v) \leq \delta |n^j|\}} d\tilde{y}_j(v)$ for all $u \in [s, t]$.

Since R is completely- S , by Lemma 7.1, and $\tilde{w}, \tilde{x}, \tilde{y}$ satisfy (i)–(iii) above, it is immediate from Theorem 5.1 of [43] that there is a constant $c_1 > 0$ depending only on R such that

$$(141) \quad \text{Osc}(\tilde{w}, [s, t]) \leq c_1 (\text{Osc}(\tilde{x}, [s, t]) + \delta \max_{j \in \mathbb{J}} |n^j|),$$

$$(142) \quad \text{Osc}(\tilde{y}, [s, t]) \leq c_1 (\text{Osc}(\tilde{x}, [s, t]) + \delta \max_{j \in \mathbb{J}} |n^j|).$$

Applying the reverse linear transformation ABA' and making \hat{C}_0 sufficiently large to absorb the factor $\max_{j \in \mathbb{J}} |n^j|$, it follows that (139)–(140) hold where the constant \hat{C}_0 can be chosen to depend only on ABA' (and its inverse). \square

Proof of Lemma 7.5. Fix $T > 0$ and $\delta > 0$. By the convergence assumed for the initial random variables $\{\hat{W}^r(0), r > 0\}$, the C -tightness of $\{\hat{W}^r(0) + \hat{X}^r(\cdot), r > 0\}$ established in Lemma 7.4, the multiplicative state space collapse established in Theorem 5.1, and the fact that $(M^r)^{-1} \rightarrow M^{-1}$ as $r \rightarrow \infty$, we have that for each $\varepsilon > 0$ there are constants $K_0 \geq 1$ (not depending on ε) and $r_0(\varepsilon) > 0$ such that for all $r \geq r_0(\varepsilon)$,

$$(143) \quad P(|\hat{W}^r(0)| \leq K_0, \|\hat{X}^r(\cdot)\|_T \leq K_0) \geq 1 - \frac{\delta}{2},$$

$$(144) \quad P(\|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T \leq \varepsilon(\|\hat{N}^r(\cdot)\|_T \vee 1)) \geq 1 - \frac{\delta}{2},$$

and

$$(145) \quad \|(M^r)^{-1} - M^{-1}\| \leq \varepsilon.$$

The constants K_0 and $r_0(\varepsilon)$ will depend on T and δ as well, but since these parameters are fixed throughout this proof, we do not explicitly indicate that dependence here.

In the following, $\varepsilon > 0$ will be fixed. A specific, suitably small value of ε will be chosen later (as a function of δ, T) to ensure that various inequalities hold (cf. (158)). For $r > 0$, let

$$\begin{aligned} \mathcal{O}^{r,\varepsilon} = & \{|\hat{W}^r(0)| \leq K_0, \|\hat{X}^r(\cdot)\|_T \leq K_0, \\ & \|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T \leq \varepsilon(\|\hat{N}^r(\cdot)\|_T \vee 1)\}. \end{aligned}$$

From (143)–(144) we have that for all $r \geq r_0(\varepsilon)$,

$$(146) \quad P(\mathcal{O}^{r,\varepsilon}) \geq 1 - \delta.$$

Now, for $r \geq r_0(\varepsilon)$, on $\mathcal{O}^{r,\varepsilon}$, by (130)–(132) and (145) we have

$$(147) \quad \|\hat{\xi}^r(\cdot)\|_T \leq \hat{C}_3 \varepsilon (\|\hat{N}^r(\cdot)\|_T \vee 1),$$

where $\hat{C}_3 = \|AM^{-1}\| + \|A\| \vee 1$ and

$$(148) \quad \hat{\xi}^r(t) = \hat{W}^r(t) - \tilde{W}^r(t) \quad \text{for all } t \geq 0.$$

Since, by Proposition 4.2, $\Delta(\hat{W}^r(t, \omega))$ is an invariant state, it follows from the characterization of invariant states in Theorem 4.1 that for $r \geq r_0(\varepsilon)$, $t \in [0, T]$ and $\omega \in \mathcal{O}^{r,\varepsilon}$, there is $q^r(t, \omega) \in \mathbb{R}_+^{\mathbb{J}}$ such that

$$(149) \quad \Delta(\hat{W}^r(t, \omega)) = \text{diag}(\rho) \text{diag}(\kappa)^{-1} A' q^r(t, \omega),$$

and so by the definitions of \tilde{W}^r and the matrix B we have

$$(150) \quad \tilde{W}^r(t, \omega) = ABA' q^r(t, \omega).$$

We now turn to the behavior of \hat{U}^r . By (39), for each $r \geq r_0(\varepsilon)$, $t \in [0, T]$, and $j \in \mathbb{J}$,

$$(151) \quad \hat{U}_j^r(t) = r \int_0^t (C_j - (A\Lambda(\hat{N}^r(s)))_j) ds.$$

We are interested in where \hat{U}_j^r increases. Fix $r \geq r_0(\varepsilon)$, $j \in \mathbb{J}$ and $\omega \in \mathcal{O}^{r,\varepsilon}$. Since the integrand in (151) is non-negative, we concentrate on determining where the integrand is strictly positive. At an instant $s \in [0, T]$ such that $C_j - (A\Lambda(\hat{N}^r(s, \omega)))_j > 0$, by Proposition 2.1(iv) there is $p^r(s, \omega) \in \mathbb{R}_+^{\mathbb{J}}$ such that $p_j^r(s, \omega) = 0$ and

$$(152) \quad \hat{N}_i^r(s, \omega) = \Lambda_i(\hat{N}^r(s, \omega)) \left(\frac{\sum_{k \in \mathbb{J}} p_k^r(s, \omega) A_{ki}}{\kappa_i} \right) \quad \text{for all } i \in \mathbb{I}.$$

By the local traffic Assumption 5.1, there is an index $i_j \in \mathbb{I}$ such that $A_{ji_j} > 0$ and $A_{ki_j} = 0$ for all $k \neq j$. Using the fact that $p_j^r(s, \omega) = 0$ in (152), it follows that $\hat{N}_{i_j}^r(s, \omega) = 0$. Then, since $\omega \in \mathcal{O}^{r,\varepsilon}$, we have

$$(153) \quad (\Delta(\hat{W}^r(s, \omega)))_{i_j} \leq \varepsilon (\|\hat{N}^r(\cdot, \omega)\|_T \vee 1)$$

and so by (149) and the local traffic condition,

$$(154) \quad q_j^r(s, \omega) \leq \varepsilon (\|\hat{N}^r(\cdot, \omega)\|_T \vee 1) \rho_{i_j}^{-1} \kappa_{i_j}.$$

Thus, letting n^j denote the vector given by the j^{th} row of $(ABA')^{-1}$ (cf. the proof of Lemma 7.2), using (150) we have that

$$\begin{aligned} d(\tilde{W}^r(s, \omega), \mathscr{W}_1^j) &= n^j \cdot \tilde{W}^r(s, \omega) / |n^j| \\ &= n^j \cdot ABA'q^r(s, \omega) / |n^j| \\ &= q_j^r(s, \omega) / |n^j| \\ &\leq \varepsilon(\|\hat{N}^r(\cdot, \omega)\|_T \vee 1) \rho_{i_j}^{-1} \kappa_{i_j} / |n^j|. \end{aligned}$$

Let $\hat{C}_4 = \max_{j \in \mathbb{J}} \rho_{i_j}^{-1} \kappa_{i_j} / |n^j|$. It follows from the reasoning above that for each $r \geq r_0(\varepsilon)$, $j \in \mathbb{J}$, on $O^{r, \varepsilon}$ we have for all $t \in [0, T]$,

$$(155) \quad \hat{U}_j^r(t) = \int_0^t \mathbf{1}_{\{d(\tilde{W}^r(s), \mathscr{W}_1^j) \leq \hat{C}_4 \varepsilon (\|\hat{N}^r(\cdot)\|_T \vee 1)\}} d\hat{U}_j^r(s),$$

i.e., $\hat{U}_j^r(\cdot)$ can increase only when $d(\tilde{W}^r(\cdot), \mathscr{W}_1^j) \leq \hat{C}_4 \varepsilon (\|\hat{N}^r(\cdot)\|_T \vee 1)$. Furthermore,

$$\begin{aligned} \tilde{W}^r(t) &= \hat{W}^r(t) - \hat{\xi}^r(t) \\ &= \hat{W}^r(0) + \hat{X}^r(t) - \hat{\xi}^r(t) + \hat{U}^r(t), \end{aligned}$$

where $\tilde{W}^r(t) \in \mathscr{W}_1$. It then follows from the oscillation inequality in Proposition 7.1 that for $r \geq r_0(\varepsilon)$, on $O^{r, \varepsilon}$,

$$\begin{aligned} \text{Osc}(\tilde{W}^r, [0, T]) &\leq \hat{C}_0 \left(\text{Osc}(\hat{W}^r(0) + \hat{X}^r(\cdot) - \hat{\xi}^r(\cdot), [0, T]) \right. \\ &\quad \left. + \hat{C}_4 \varepsilon (\|\hat{N}^r(\cdot)\|_T \vee 1) \right) \\ (156) \quad &\leq \hat{C}_0 (2K_0 + (2\hat{C}_3 + \hat{C}_4) \varepsilon (\|\hat{N}^r(\cdot)\|_T \vee 1)) \\ &\leq \hat{C}_0 (2K_0 + (2\hat{C}_3 + \hat{C}_4) \varepsilon \hat{C}_2 (\|\hat{W}^r(\cdot)\|_T \vee 1)). \end{aligned}$$

where we have used the definition of $O^{r, \varepsilon}$ and (147) for the second inequality, and we have used Lemma 7.3 plus the fact that $\hat{C}_2 \geq 1$ for the third inequality. Similarly, we obtain an oscillation bound for \hat{U}^r for $r \geq r_0(\varepsilon)$ on $O^{r, \varepsilon}$:

$$(157) \quad \text{Osc}(\hat{U}^r, [0, T]) \leq \hat{C}_0 (2K_0 + (2\hat{C}_3 + \hat{C}_4) \varepsilon \hat{C}_2 (\|\hat{W}^r(\cdot)\|_T \vee 1)).$$

On combining (156) with (148), (147) and Lemma 7.3, we have that for $r \geq r_0(\varepsilon)$, on $O^{r, \varepsilon}$,

$$\begin{aligned} \|\hat{W}^r(\cdot)\|_T \vee 1 &\leq |\hat{W}^r(0)| \vee 1 + \text{Osc}(\hat{W}^r(\cdot), [0, T]) \\ &\leq |\hat{W}^r(0)| \vee 1 + \text{Osc}(\tilde{W}^r(\cdot), [0, T]) + \text{Osc}(\hat{\xi}^r(\cdot), [0, T]) \\ &\leq K_0 + (\hat{C}_0 + 1) (2K_0 + (2\hat{C}_3 + \hat{C}_4) \varepsilon \hat{C}_2 (\|\hat{W}^r(\cdot)\|_T \vee 1)). \end{aligned}$$

Now, choose

$$(158) \quad \varepsilon = \min \left(\frac{1}{2(\hat{C}_0 + 1)(2\hat{C}_3 + \hat{C}_4)\hat{C}_2}, \frac{\delta}{2K_0\hat{C}_2(3 + 2\hat{C}_0)(\hat{C}_3 + \hat{C}_4)} \right).$$

Then from the above we conclude that for $r \geq r_0(\varepsilon)$, on $\mathcal{O}^{r,\varepsilon}$,

$$(159) \quad \|\hat{W}^r(\cdot)\|_T \vee 1 \leq 2K_0(3 + 2\hat{C}_0),$$

and, using the fact that $\hat{C}_3 \geq 1$, we have

$$(160) \quad \begin{aligned} \varepsilon(\|\hat{N}^r(\cdot)\|_T \vee 1) &\leq \varepsilon\hat{C}_2(\|\hat{W}^r(\cdot)\|_T \vee 1) \\ &\leq \frac{\delta}{\hat{C}_3 + \hat{C}_4} \\ &\leq \delta. \end{aligned}$$

Further, by (157), (159) and the definition of ε , we have

$$(161) \quad \text{Osc}(\hat{U}^r, [0, T]) \leq 2\hat{C}_0(K_0 + \delta).$$

Note that since \hat{U}^r is non-decreasing and starts from zero, the above also provides a bound for $\|\hat{U}^r(\cdot)\|_T$. Now, by (160), for $r \geq r_0(\varepsilon)$, on $\mathcal{O}^{r,\varepsilon}$, we have

$$\begin{aligned} \|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T &\leq \varepsilon(\|\hat{N}^r(\cdot)\|_T \vee 1) \\ &\leq \delta, \end{aligned}$$

and using (147), we have

$$\begin{aligned} \|\hat{\xi}^r(\cdot)\|_T &\leq \hat{C}_3\varepsilon(\|\hat{N}^r(\cdot)\|_T \vee 1) \\ &\leq \delta, \end{aligned}$$

and using (155), we have for each $j \in \mathbb{J}$ and $t \in [0, T]$,

$$(162) \quad \hat{U}_j^r(t) = \int_0^t 1_{\{d(\hat{W}^r(s, \omega), \mathcal{N}_1^j) \leq \delta\}} d\hat{U}_j^r(s),$$

and so by (134)–(135),

$$(163) \quad \|\hat{\xi}^{r,\delta}(\cdot)\|_T = 0.$$

Now, set $K(T, \delta) = \max(2K_0(3 + 2\hat{C}_0), 2\hat{C}_0(K_0 + \delta))$ and $r(T, \delta) = r_0(\varepsilon)$ where $\varepsilon > 0$ is given by (158). Then, for all $r \geq r(T, \delta)$, by the above and (146) we have

$$(164) \quad \begin{aligned} P \left(\|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_T \leq \delta, \|\hat{\xi}^r(\cdot)\|_T \leq \delta, \|\hat{\xi}^{r,\delta}(\cdot)\|_T = 0, \right. \\ \left. \|\hat{W}^r(\cdot)\|_T \leq K(T, \delta), \|\hat{U}^r(\cdot)\|_T \leq K(T, \delta) \right) \geq P(\mathcal{O}^{r,\varepsilon}) \geq 1 - \delta. \end{aligned}$$

□

Proof of Theorem 5.2. Assume that the hypotheses of Theorem 5.2 hold. We will show that the conditions of Theorem 5.4 of [20] hold, from which it will follow immediately that \hat{W}^r converges in distribution as $r \rightarrow \infty$ to an SRBM with data $(\mathscr{W}_1, \theta, \Gamma, \{\gamma^j : j \in \mathbb{J}\}, \eta)$. The joint convergence of \hat{N}^r with \hat{W}^r will then follow by the state space collapse established in Lemma 7.5.

The conditions of Theorem 5.4 fall into four groups. We treat each of these groups separately below.

Firstly, the diffusion state space \mathscr{W}_1 must be a convex polyhedron having non-empty interior described as the intersection of a minimal set of half-spaces, and the directions of reflection $\{\gamma^j : j \in \mathbb{J}\}$ must satisfy Assumption 5.1 of [20]. These properties were verified in the proof of Lemma 7.2.

Secondly, one must verify that Assumption 4.1 of [20] holds. This condition amounts to verifying that $\hat{W}^r, \hat{X}^r, \hat{U}^r$ satisfy conditions similar to (i), (ii), (iv) on $\tilde{W}, \tilde{X}, \tilde{U}$ in the definition of an SRBM, except that the prelimit processes are allowed to have r.c.l.l. rather than continuous paths and small perturbations in the conditions (i), (ii) and (iv) are allowed. Furthermore, the sequence $\{\hat{W}^r(0) + \hat{X}^r(\cdot)\}$ is required to be C -tight. The latter follows from Lemma 7.4 and the former follows from the properties in (129)–(135) once we show that $\hat{\xi}^r \rightarrow 0$ and $\hat{\zeta}^{r, \delta^r} \rightarrow 0$ in probability as $r \rightarrow \infty$ for a suitable sequence $\{\delta^r\}$ satisfying $\delta^r \rightarrow 0$ as $r \rightarrow \infty$. We verify the latter properties in the next paragraph.

Recall the constants $K(T, \delta)$ and $r(T, \delta)$ from Lemma 7.5. Choose strictly increasing sequences of positive constants $\{K_k, k \geq 1\}$ and $\{r_k, k \geq 1\}$ such that for each k , $K_k \geq K(k, \frac{1}{k})$ and $r_k \geq r(k, \frac{1}{k})$, and so that $r_k \rightarrow \infty$ as $k \rightarrow \infty$. Define δ^r such that $\delta^r = 1$ when $r \leq r_1$ and $\delta^r = \frac{1}{k}$ when $r \in (r_k, r_{k+1}]$ for $k \geq 1$. Then, $\delta^r \rightarrow 0$ as $r \rightarrow \infty$ and by Lemma 7.5, for each $k \geq 1$, for $r_k < r \leq r_{k+1}$,

$$(165) \quad \begin{aligned} P\left(\|\hat{N}^r(\cdot) - \Delta(\hat{W}^r(\cdot))\|_k \leq \frac{1}{k}, \|\hat{\xi}^r(\cdot)\|_k \leq \frac{1}{k}, \|\hat{\zeta}^{r, \delta^r}(\cdot)\|_k = 0, \right. \\ \left. \|\hat{W}^r(\cdot)\|_k \leq K_k, \|\hat{U}^r(\cdot)\|_k \leq K_k\right) \geq 1 - \frac{1}{k}. \end{aligned}$$

It follows from this that $\hat{\xi}^r \rightarrow 0$ and $\hat{\zeta}^{r, \delta^r} \rightarrow 0$ in probability as $r \rightarrow \infty$. Then the conditions of Assumption 4.1 in [20] are satisfied with \mathscr{W}_1 in place of \bar{G} , r in place of n , j in place of i , \hat{W}^r in place of W^n , \hat{W}^r in place of \tilde{W}^n , $\hat{W}^r(0) + \hat{X}^r$ in place of X^n , \hat{U}^r in place of Y^n , $\tilde{Y}_j^n(t) = \int_0^t 1_{\{d(\tilde{W}^r(s), \mathscr{W}_1^j) \leq \delta\}} d\hat{U}_j^r(s)$, $\gamma^{j, n} = \gamma^j$, $\alpha^n = \hat{\xi}^r$, and $\beta^n = \hat{\zeta}^{r, \delta^r}$.

Thirdly, one must show that $\{\hat{W}^r(0) + \hat{X}^r(\cdot)\}$ converges in distribution as $r \rightarrow \infty$ to a Brownian motion with drift θ , covariance matrix Γ given by (46) and initial distribution η . [This amounts to verifying condition (vi)' of Theorem 4.3 in [20] (with θ in place of μ and η in place of ν there). This condition is needed so that weak limit points of $\{(\hat{W}^r, \hat{X}^r, \hat{U}^r)\}$ will satisfy property (iii)(a) and the last part of

(ii) in the Definition 5.2 of an SRBM.] To verify this condition, we first show that $\{\bar{T}^r(\cdot)\}$ converges in distribution to the deterministic process $\{\rho(t), t \geq 0\}$ where $\rho(t) = \rho t$. Indeed, by the fact that $\bar{T}^r(\cdot)$ is uniformly Lipschitz continuous with a Lipschitz constant which does not depend on r , we have that $\{\bar{T}^r(\cdot), r > 0\}$ is C -tight. Let $T^*(\cdot)$ be a weak limit point of this sequence, obtained as a limit in distribution along a suitable subsequence. We will prove below that almost surely,

$$(166) \quad T^*(t) = \rho t \quad \text{for all } t \geq 0,$$

from which the desired convergence in distribution of $\{\bar{T}^r(\cdot)\}$ to $T^*(\cdot)$ follows.

For the proof of (166), by passing to a subsequence, we may assume that $\{\bar{T}^r(\cdot)\}$ converges in distribution to $T^*(\cdot)$ as $r \rightarrow \infty$. Now, by (16) and (5), on dividing by r , we have for each $i \in \mathbb{I}$,

$$(167) \quad \frac{\hat{N}_i^r(t)}{r} = \frac{\hat{N}_i^r(0)}{r} + \bar{E}_i^r(t) - \bar{S}_i^r(\bar{T}_i^r(t)),$$

where

$$(168) \quad \bar{E}_i^r(t) = \frac{E^r(r^2 t)}{r^2}, \quad \bar{S}_i^r(t) = \frac{S^r(r^2 t)}{r^2}.$$

From (165) we can conclude using the continuity of $\Delta(\cdot)$ that

$$\frac{\hat{N}_i^r(\cdot)}{r} \Rightarrow 0 \quad \text{as } r \rightarrow \infty,$$

and from (21) it follows that

$$(\bar{E}_i^r(\cdot), \bar{S}_i^r(\cdot)) \Rightarrow (v(\cdot), \mu(\cdot)) \quad \text{as } r \rightarrow \infty,$$

where $v(t) = vt$ and $\mu(t) = \mu t$ for all $t \geq 0$. Thus, on letting $r \rightarrow \infty$ in (167), we obtain that almost surely, for each $i \in \mathbb{I}$,

$$0 = v_i t - \mu_i T_i^*(t) \quad \text{for all } t \geq 0.$$

The desired result (166) follows immediately.

It now follows from (21) and the assumption on the convergence in distribution of $\hat{W}^r(0)$ made at the end of Section 3, that $(\hat{W}^r(0), \hat{E}^r(\cdot), \hat{S}^r(\cdot), \bar{T}^r(\cdot))$ converges in distribution to $(\tilde{W}(0), \tilde{E}(\cdot), \tilde{S}(\cdot), \rho(\cdot))$ as $r \rightarrow \infty$ where $\tilde{W}(0)$ is independent of the Brownian motion $(\tilde{E}(\cdot), \tilde{S}(\cdot))$ and $\tilde{W}(0)$ has distribution η on \mathscr{W}_1 . On combining this with a random time change theorem and Assumption 3.1, we conclude using (128) that $\{\hat{W}^r(0) + \hat{X}^r(\cdot), r > 0\}$ converges in distribution to a Brownian motion with drift θ , covariance matrix Γ given by (46), and initial distribution η , as desired.

Fourthly, and finally, we must verify condition (vii) of Theorem 4.3 of [20] (with θ in place of μ there). This condition requires that for any weak limit point (W, X, U) of $(\hat{W}^r, \hat{X}^r, \hat{U}^r)$, $\{X(t) - \theta t, t \geq 0\}$ is a martingale relative to the filtration generated by (W, X, U) . (This condition is needed so that property (iii)(b) in Definition 5.2 of an SRBM will be satisfied by weak limit points of $\{(\hat{W}^r, \hat{X}^r, \hat{U}^r)\}$.) Let $\hat{\theta}^r = r(A\rho^r - C)$. Then, by Proposition 4.1 of [20], to show that this final condition holds, it suffices to verify the following properties for the prelimit processes:

- (a) for each r , $\{\hat{X}^r(t) - \hat{\theta}^r t, t \geq 0\}$ is a martingale with respect to the filtration generated by $(\hat{W}^r, \hat{X}^r, \hat{U}^r)$,
- (b) $\hat{\theta}^r \rightarrow \theta$ as $r \rightarrow \infty$,
- (c) $\{\hat{X}^r(t), r > 0\}$ is uniformly integrable for each fixed $t \geq 0$.

We now verify these three properties.

It is well known (cf. Theorem 4.1, Chapter 6 of [12]) that for the continuous time Markov chain N^r , for each $r > 0$ and $i \in \mathbb{I}$,

$$\chi_i^r(t) = N_i^r(t) - N_i^r(0) - \int_0^t (v_i^r - \mu_i^r \Lambda_i(N^r(s))) ds, \quad t \geq 0,$$

defines a martingale with respect to the filtration generated by N^r , i.e., with respect to $\{\mathcal{F}_t^r, t \geq 0\}$ where $\mathcal{F}_t^r = \sigma\{N^r(s) : 0 \leq s \leq t\}$. Setting

$$\hat{\chi}^r(t) = \frac{\chi^r(r^2 t)}{r}, \quad \hat{\mathcal{F}}_t^r = \mathcal{F}_{r^2 t}^r,$$

we have that for each $r > 0$, $\{\hat{\chi}^r(t), \hat{\mathcal{F}}_t^r, t \geq 0\}$ is a multi-dimensional martingale. Now, using the expressions (17) and (36) for \hat{W}^r , and the expression (39) for \hat{U}^r , we see that

$$\begin{aligned} A(M^r)^{-1} \hat{\chi}^r(t) &= \hat{W}^r(t) - \hat{W}^r(0) - A\rho^r r t + C r t - \hat{U}^r(t) \\ &= \hat{X}^r(t) + r(C - A\rho^r)t. \end{aligned}$$

Since \hat{W}^r and \hat{U}^r are adapted to the filtration $\{\hat{\mathcal{F}}_t^r, t \geq 0\}$, so is \hat{X}^r , and then it follows that

$$\hat{X}^r(t) - \hat{\theta}^r t = A(M^r)^{-1} \hat{\chi}^r(t), \quad t \geq 0,$$

is a martingale with respect to the filtration generated by $(\hat{W}^r, \hat{X}^r, \hat{U}^r)$. By (12), $\hat{\theta}^r \rightarrow \theta$ as $r \rightarrow \infty$. Finally, since the mean of a Poisson random variable is the same as its variance, and the rate of increase of \bar{T}^r is uniformly bounded, it is straightforward to verify using Assumption 3.1 that $\{\hat{X}^r(t), r > 0\}$ is uniformly integrable for each $t \geq 0$. Thus we have verified conditions (a)–(c) above and so condition (vii) of Theorem 4.3 in [20] holds.

We have now verified all of the hypotheses of Theorem 5.4 of [20] and so it follows that \hat{W}^r converges in distribution as $r \rightarrow \infty$ to an SRBM \tilde{W} associated with the data $(\mathcal{W}_1, \theta, \Gamma, \{\gamma^j : j \in \mathbb{J}\}, \eta)$. From (165) we have that $\hat{N}^r - \Delta(\hat{W}^r)$ converges in distribution to the zero process as $r \rightarrow \infty$. Since $\Delta(\cdot)$ is continuous, it follows that we have the joint convergence of (\hat{W}^r, \hat{N}^r) in distribution as $r \rightarrow \infty$ to $(\tilde{W}, \Delta(\tilde{W}))$. \square

8. CONCLUSION. A bandwidth sharing policy corresponds to a generalization of the notion of a processor sharing discipline from a single resource to a network with several shared resources. In particular, weighted α -fair policies provide a tractable theoretical abstraction of the bandwidth sharing effected by decentralized packet-based end-to-end congestion control algorithms such as TCP. It is known [3, 11] that, for flow-level models with exponentially distributed file sizes, weighted α -fair policies are stable when the average load on each resource is less than its capacity.

Weighted α -fair policies can nevertheless suffer from entrainment of resources, whereby congestion at some resources may prevent others from working at full capacity: this is manifested under diffusion scaling, where a Brownian model for the workload process lives in a cone which may be a strict subset of the positive orthant.

Under weighted proportional fair sharing ($\alpha = 1$) and a mild local traffic condition, this paper has shown how multiplicative state space collapse can be combined with an invariance principle to establish a diffusion approximation for the flow-level model. For proportional fair sharing (equal weights), this diffusion has a product form invariant measure which when integrable can be normalized to yield the unique stationary distribution for the diffusion. This result extends to the case where file sizes are distributed as finite mixtures of exponentials. Thus, in the diffusion limit more networks can have product form stationary distributions than are known in the prelimit (cf. Bonald and Proutière [4]). We also indicated some open problems for $\alpha \neq 1$ and an extension to models with routing.

Bandwidth sharing policies outside the class of weighted α -fair policies may avoid entrainment, although such policies may not be easy to effect via decentralized packet-based end-to-end congestion control algorithms.

REFERENCES

- [1] Ben Fredj, S., Bonald, T., Proutière, A., Regnie, G. and Roberts, J. (2001). Statistical bandwidth sharing: a study of congestion at flow level. *Computer Communication Review* **31:4** 111–121.
- [2] Billingsley, P. (1999). *Convergence of Probability Measures*. Second edition, John Wiley & Sons Inc., New York.
- [3] Bonald, T. and Massoulié, L. (2001). Impact of fairness on Internet performance. *Proceedings of ACM Sigmetrics 2001*.

- [4] Bonald T. and Proutière A. (2003). Insensitive bandwidth sharing in data networks. *Queueing Systems* **44** 69–100.
- [5] Bramson, M. (1998). State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems: Theory and Applications* **30** 89–148.
- [6] Bramson, M. (2005). Stability of networks for max-min fair routing. Presentation at the 13th INFORMS Applied Probability Conference, Ottawa.
- [7] Bu, T. and Towsley, D. (2001). Fixed point approximation for TCP behavior in an AQM network. *Proceedings of ACM Sigmetrics 2001*.
- [8] Chiang, M., Shah, D. and Tang, A. K. (2006). Stochastic stability under network utility maximization: general file size distribution. In *Proceedings of 44th Allerton Conference on Communication, Control and Computing*.
- [9] Cottle, R. W., Pang, J.-S. and Stone, R. E. (1992). *The Linear Complementarity Problem*. Academic Press, San Diego.
- [10] Dai, J. G. and Williams, R. J. (1995). Existence and uniqueness of semimartingale reflecting Brownian motions in convex polyhedrons, *Theory of Probability and Its Applications* **40** 1–40. Correction: **50** (2006), 346–347.
- [11] De Veciana, G., Lee, T. J. and Konstantopoulos, T. (2001). Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Transactions on Networking* **9** 2–14.
- [12] Ethier, S. N. and Kurtz, T. G. *Markov Processes: Characterization and Convergence*. Wiley, New York, 1986.
- [13] Gibbens, R. J., Sargood, S. K., Van Eijl, C., Kelly, F. P., Azmoodeh, H., Macfadyen, R. N. and Macfadyen, N. W. (2000). Fixed-point models for the end-to-end performance analysis of IP networks. *13th ITC Specialist Seminar: IP Traffic Measurement, Modeling and Management*, Sept 2000, Monterey, California
- [14] Gromoll, H. C. and Williams, R. J. (2008). Fluid limits for networks with bandwidth sharing and general document size distributions, to appear in *Annals of Applied Probability*.
- [15] Gromoll, H. C. and Williams, R. J. (2007). Fluid model for a data network with α -fair bandwidth sharing and general document size distributions: two examples of stability. To appear in *Proceedings of Markov Processes and Related Topics* — a conference in honor of Tom Kurtz's 65th birthday.
- [16] Harrison, J. M. (2000). Brownian models of open processing networks: canonical representation of workload. *Annals of Applied Probability* **10** 75–103. Correction: **13** (2003), 390–393.
- [17] Harrison, J. M. and Williams, R. J. (1987). Multidimensional reflected Brownian motions having exponential stationary distributions. *Annals of Probability* **15** 115–137.
- [18] Harrison, J. M. and Williams, R. J. (1987). Brownian models of open queueing networks with homogeneous customer populations. *Stochastics* **22** 77–115.
- [19] Han, H., Shakkottai, S., Hollot, C.V., Srikant, R. and Towsley, D. (2006). Multi-path TCP: a joint congestion control and routing scheme to exploit path diversity on the Internet. *IEEE/ACM Transactions on Networking* **14** 1260–1271.
- [20] Kang, W. and Williams, R. J. (2007). An invariance principle for semimartingale reflecting Brownian motions in domains with piecewise smooth boundaries. *Annals of Applied Probability*, **17**, 741–779.
- [21] Kang, W. N. and Williams, R. J. (2007). Diffusion approximation for an input-queued packet switch operating under a maximum weight algorithm. In preparation.
- [22] Kang, W., Kelly, F. P., Lee, N. H. and Williams, R. J. (2004). On fluid and Brownian approximations for an Internet congestion control model. *Proceedings of the 43rd IEEE Conference on Decision and Control*. December 2004, 3938–3943.
- [23] Kang, W. N., Kelly, F. P., Lee, N. H. and Williams, R. J. (2007). Product form stationary distributions for diffusion approximations to a flow-level model operating under a proportional fair sharing policy, *ACM SIGMETRICS Performance Evaluation Review*, Vol. 35, Issue 2 (2007),

- 36–38.
- [24] Kelly, F. P. (1991). Loss networks. *Annals of Applied Probability* **1** 319–378.
 - [25] Kelly, F. P. and Voice, T. (2005). Stability of end-to-end algorithms for joint routing and rate control. *Computer Communication Review* **35** :2, 5–12.
 - [26] Kelly, F. P. and Williams, R. J. (2004). Fluid model for a network operating under a fair bandwidth-sharing policy. *Annals of Applied Probability* **14** 1055–1083.
 - [27] Key, P. and Massoulié, L. (2006). Fluid models of integrated traffic and multipath routing. *Queueing Systems* **53** 85–98.
 - [28] Lakshminantha, A., Beck, C. L. and Srikant, R. (2004). Connection level stability analysis of the Internet using the sum of squares (SoS) techniques. *Conference on Information Sciences and Systems, Princeton*.
 - [29] Laws, C. N. (1992). Resource pooling in queueing networks with dynamic routing. *Advances in Applied Probability* **24** 699–726.
 - [30] Lin, X. and Shroff, N. (2004). On the stability region of congestion control. *Proceedings of the Allerton Conference on Communications, Control and Computing*.
 - [31] Lin, X., Shroff, N. and Srikant, R. (2006). On the connection-level stability of congestion-controlled communication networks. *IEEE Transactions on Information Theory*, to appear.
 - [32] Massoulié, L. (2007). Structural properties of proportional fairness: stability and insensitivity. *Annals of Applied Probability*, **17**, 809–839.
 - [33] Massoulié, L. and Roberts, J. (2000). Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, **15**, 185–201.
 - [34] Massoulié, L. and Roberts, J. (2002). Bandwidth Sharing: Objectives and Algorithms. *IEEE/ACM Transactions on Networking*, **10**, 320–328.
 - [35] Mo, J. and Walrand, J. (2000). Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking* **8** 556–567.
 - [36] Roughan, M., Erramilli, A. and Veitch, D. (2001). Network performance for TCP networks, Part I: persistent sources. *Proceeding ITC'17*, Brasil, September 2001.
 - [37] Schassberger, R. (1986). Two remarks on insensitive stochastic models. *Advances in Applied Probability* **18** 791–814.
 - [38] Shah, D. and Wischik, D. (2007). Heavy traffic analysis of optimal scheduling algorithms for switched networks. Preprint.
 - [39] Srikant, R. (2004). On the positive recurrence of a Markov chain describing file arrivals and departures in a congestion-controlled network. *Presented at the IEEE Computer Communications Workshop*.
 - [40] Tassiulas, L. and Ephremides, A. (1992). Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks. *IEEE Transactions on Automatic Control* **37** 1936–1948.
 - [41] Williams, R. J. (1987). Reflected Brownian motion with skew symmetric data in a polyhedral domain. *Probability Theory and Related Fields* **75** 459–485.
 - [42] Williams, R. J. (1998). Diffusion approximations for open multiclass queueing networks: sufficient conditions involving state space collapse. *Queueing Systems: Theory and Applications* **30** 27–88.
 - [43] Williams, R. J. (1998). An invariance principle for semimartingale reflecting Brownian motions in an orthant. *Queueing Systems*, **30** 5–25.
 - [44] Ye, H. and Yao, D. D. (2007). Heavy-traffic optimality of a stochastic network under utility-maximizing resource allocation, to appear in *Operations Research*.

DEPARTMENT OF MATHEMATICAL SCIENCES
CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PA 15213-3890
USA
E-MAIL: weikang@andrew.cmu.edu

STATISTICAL LABORATORY
CENTER FOR MATHEMATICAL SCIENCES
UNIVERSITY OF CAMBRIDGE
WILBERFORCE ROAD
CAMBRIDGE CB3 0WB
UNITED KINGDOM
E-MAIL: F.P.Kelly@statslab.cam.ac.uk

DEPARTMENT OF MATHEMATICS
UNIVERSITY OF CALIFORNIA, SAN DIEGO,
9500 GILMAN DRIVE
LA JOLLA, CA 92093-0112
USA
E-MAIL: nhlee@math.ucsd.edu; williams@math.ucsd.edu