

Product Form Stationary Distributions for Diffusion Approximations to a Flow-Level Model Operating under a Proportional Fair Sharing Policy

W. N. Kang
 Department of Mathematical Sciences
 Carnegie Mellon University
 Pittsburgh PA 15213-3890
 USA

F. P. Kelly*
 Center for Mathematical Sciences
 University of Cambridge
 Cambridge CB3 0WB
 United Kingdom

N. H. Lee and R. J. Williams†
 Department of Mathematics
 University of California, San Diego
 La Jolla CA 92093-0112
 USA

1. INTRODUCTION

We consider a flow-level model of Internet congestion control introduced by Massoulié and Roberts [2]. We assume that bandwidth is shared amongst elastic documents according to a weighted proportional fair bandwidth sharing policy. With 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. In [1], under a mild local traffic condition, we establish a diffusion approximation for the workload process (and hence for the flow count process) in this model. We first recall that result in this paper. We then state results showing that when all of the weights are equal (proportional fair sharing) the diffusion has a product form invariant distribution with a strikingly simple interpretation in terms of dual random variables, one for each of the resources of the network. This result can be extended to the case where document sizes are distributed as finite mixtures of exponentials, and to models that include multi-path routing (these extensions are not described here, but can be found in [1]).

2. STOCHASTIC FLOW-LEVEL MODEL

Consider a network with finitely many resources labeled by $j \in \mathcal{J}$ and finitely many routes labeled by $i \in \mathcal{I}$. A route is a non-empty subset of the resources (interpreted as the set of resources used simultaneously by a flow on that route). Let $\mathbf{J} = |\mathcal{J}|$, the total number of resources, and $\mathbf{I} = |\mathcal{I}|$, the total number of routes. Set $A_{ji} = 1$ if resource j is used by route i , and set $A_{ji} = 0$ otherwise. This defines a $\mathbf{J} \times \mathbf{I}$ matrix A that is assumed to have full row rank. It is assumed that resource capacities ($C_j : j \in \mathcal{J}$) are given and that these are all strictly positive and finite.

A 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 a given route. The flow count process is assumed to be a Markov process with state space $\mathbb{N}^{\mathbf{I}}$, where \mathbb{N} denotes the set of non-negative in-

tegers. In particular, 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. Bandwidth capacity is allocated dynamically to the documents according to the following bandwidth sharing policy which is called *weighted proportional fair sharing*. Given strictly positive weights ($\kappa_i : i \in \mathcal{I}$), if $N_i(t)$ denotes the (random) number of flows on route i at time t for each $i \in \mathcal{I}$, and $N(t) = (N_i(t) : i \in \mathcal{I})$, 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 $\Lambda(\cdot) = (\Lambda_i(\cdot) : i \in \mathcal{I})$ is defined as follows. Let $\Lambda : \mathbb{R}_+^{\mathbf{I}} \rightarrow \mathbb{R}_+^{\mathbf{I}}$ be defined such that for each $n \in \mathbb{R}_+^{\mathbf{I}}$, $\Lambda(n)$ is the unique value of Λ that solves the optimization problem:

$$\begin{aligned} & \text{maximize} && \sum_{i \in \mathcal{I}_+(n)} \kappa_i n_i \log \Lambda_i \\ & \text{subject to} && A\Lambda \leq C \\ & \text{over} && \Lambda \in \mathcal{O}_+(n), \end{aligned} \quad (1)$$

where $\mathcal{I}_+(n) = \{i \in \mathcal{I} : n_i > 0\}$ and $\mathcal{O}_+(n) = \{\Lambda \in \mathbb{R}_+^{\mathbf{I}} : \Lambda_i = 0 \text{ for all } i \text{ such that } n_i = 0\}$. When the weights ($\kappa_i, i \in \mathcal{I}$) are all equal, the bandwidth sharing policy is simply called *proportional fair sharing*.

We define an (average) *workload process* by $W(t) = AM^{-1}N(t)$ 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. HEAVY TRAFFIC AND SCALING

Consider a sequence of flow-level models indexed by r (tending to infinity through a sequence) where the network structure with parameters A and C and bandwidth sharing policy with parameters ($\kappa_i, i \in \mathcal{I}$) 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 . Let

*Research supported in part by EPSRC Grant GR/586266/01.

†Research supported in part by NSF Grant DMS-0604537.

$\rho_i^r = v_i^r / \mu_i^r$ for each $i \in \mathcal{I}$. 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 $i \in \mathcal{I}$,*

$$v^r \rightarrow v \text{ and } \mu^r \rightarrow \mu \text{ as } r \rightarrow \infty, \quad (2)$$

$$r(A\rho^r - C) \rightarrow \theta \text{ as } r \rightarrow \infty. \quad (3)$$

Let $M = \text{diag}(\mu)$ and $\rho_i = \frac{v_i}{\mu_i}$ for all $i \in \mathcal{I}$. We note that (2)–(3) imply that $\rho^r \rightarrow \rho$ as $r \rightarrow \infty$ and $A\rho = C$.

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

$$\hat{N}^r(t) = \frac{N^r(r^2t)}{r}, \quad (4)$$

$$\hat{W}^r(t) = \frac{W^r(r^2t)}{r} = A(M^r)^{-1}\hat{N}^r(t). \quad (5)$$

4. DIFFUSION APPROXIMATION

Define

$$\Delta = \text{diag}(\rho)\text{diag}(\kappa)^{-1}A'(ABA')^{-1}, \quad (6)$$

where A' denotes the transpose of A and B is an $\mathbf{I} \times \mathbf{I}$ diagonal matrix with the i^{th} diagonal entry being $\frac{v_i}{\mu_i^2 \kappa_i} > 0$. (The inverse of ABA' exists because A was assumed to have full row rank.)

Let

$$\mathcal{W} = \{ABA'q : q \in \mathbb{R}_+^{\mathbf{J}}\}. \quad (7)$$

Thus, \mathcal{W} is a polyhedral cone; we call it the workload cone. For each $j \in \mathcal{J}$, let

$$\mathcal{W}^j = \{ABA'q : q \in \mathbb{R}_+^{\mathbf{J}} \text{ and } q_j = 0\}. \quad (8)$$

Let Γ be the $\mathbf{J} \times \mathbf{J}$ matrix given by

$$\Gamma = 2AM^{-1}\text{diag}(v)M^{-1}A'. \quad (9)$$

For each $j \in \mathcal{J}$, let γ^j denote the unit vector that is parallel to the positive j^{th} coordinate direction in $\mathbb{R}_+^{\mathbf{J}}$. Let η be a probability distribution on \mathcal{W} , endowed with the Borel σ -algebra.

Our diffusion approximation will be a Semimartingale Reflecting Brownian Motion as defined below.

DEFINITION 4.1. *A Semimartingale Reflecting Brownian Motion (SRBM) that lives in the cone \mathcal{W} , has direction of reflection γ^j on the boundary face \mathcal{W}^j for each $j \in \mathcal{J}$, has drift θ and covariance matrix Γ , and has initial distribution η on \mathcal{W} , 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, $\tilde{W}(t) \in \mathcal{W}$ 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,*
- (v) *for each $j \in \mathcal{J}$, \tilde{U}_j is an adapted, one-dimensional process such that P-a.s.,*

$$\text{(a) } \tilde{U}_j(0) = 0,$$

$$\text{(b) } \tilde{U}_j \text{ is continuous and non-decreasing,}$$

$$\text{(c) } \tilde{U}_j(t) = \int_0^t \mathbf{1}_{\{\tilde{W}(s) \in \mathcal{W}^j\}} d\tilde{U}_j(s) \text{ for all } t \geq 0.$$

Remark: Here adapted means adapted to the filtration $\{\mathcal{F}_t\}$. We call a process satisfying the above properties an SRBM associated with the data $(\mathcal{W}, \theta, \Gamma, \{\gamma^j : j \in \mathcal{J}\}, \eta)$. For more on SRBMs, see [3].

The following can be interpreted as a *local traffic* assumption under which each resource has at least one route that only uses that resource.

ASSUMPTION 4.1. (Local Traffic) *For each $j \in \mathcal{J}$, there is $i \in \mathcal{I}$ such that $A_{ji} = 1$ and $A_{ki} = 0$ for all $k \neq j$.*

THEOREM 4.1. [1] *Assume that the local traffic Assumption 4.1 holds. Suppose that $\tilde{W}^r(0)$ converges in distribution to a random variable with distribution η concentrated on \mathcal{W} , and*

$$|\hat{N}^r(0) - \Delta\hat{W}^r(0)| \rightarrow 0 \quad (10)$$

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 $\tilde{N} = \Delta\tilde{W}$ and \tilde{W} is an SRBM with data $(\mathcal{W}, \theta, \Gamma, \{\gamma^j : j \in \mathcal{J}\}, \eta)$.

5. PRODUCT FORM INVARIANT DISTRIBUTION FOR PROPORTIONAL FAIR SHARING

Massoulié and Roberts [2] showed that for a linear network with unit capacity resources satisfying $A\rho < C$, the flow count process for the original stochastic model has a product form stationary distribution under the proportional fair sharing discipline. The following result yields a diffusion analogue of that result, but for more general network structures.

THEOREM 5.1. [1] *Suppose that $\kappa_i = 1$ for all $i \in \mathcal{I}$. Let π be the measure on \mathcal{W} that is absolutely continuous with respect to Lebesgue measure with density given by*

$$p(w) = \exp(v \cdot w), \quad w \in \mathcal{W}, \quad (11)$$

where

$$v = 2\Gamma^{-1}\theta. \quad (12)$$

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

The product form of the density (11) does not imply that, when $\theta_j < 0$ for all $j \in \mathcal{J}$, the components of the SRBM \tilde{W} are independent under the stationary distribution for the SRBM, since in general the cone \mathcal{W} is not an orthant. Independence can, however, be deduced for the components of the stationary distribution for the SRBM \tilde{Q} defined below. These components can be interpreted as dual variables associated with the optimization problem (1).

COROLLARY 5.1. [1] *Suppose that the assumptions of Theorem 4.1 hold, that $\kappa_i = 1$ for all $i \in \mathcal{I}$, and $\theta_j < 0$ for all $j \in \mathcal{J}$. Let (\tilde{W}, \tilde{N}) be the process identified in Theorem 4.1, and let*

$$\tilde{Q} = 2\Gamma^{-1}\tilde{W}. \quad (13)$$

Then \tilde{Q} is an SRBM with state space $\mathbb{R}_+^{\mathcal{J}}$. It has a unique stationary distribution, and this distribution has a density relative to Lebesgue measure that is proportional to $\exp(\theta \cdot q)$, $q \in \mathbb{R}_+^{\mathcal{J}}$. Under this stationary distribution the components of \tilde{Q} are independent and \tilde{Q}_j is exponentially distributed with parameter $-\theta_j$ for each $j \in \mathcal{J}$.

It follows that the stationary distribution of $\tilde{N} = \text{diag}(\rho)A'\tilde{Q}$ 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 $j \in \mathcal{J}$; these dual variables are independent and exponentially distributed with parameters $-\theta_j, j \in \mathcal{J}$; and under its stationary distribution the i^{th} component of \tilde{N} is proportional (at any time) 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

$$N_i^s \approx \rho_i \sum_{j \in \mathcal{J}} Q_j^s A_{ji} \quad (14)$$

where $Q_j^s, j \in \mathcal{J}$, are independent and Q_j^s is exponentially distributed with parameter $C_j - \sum_{i \in \mathcal{J}} A_{ji} \rho_i$. This formal approximation accords well with the exact results of Massoulié and Roberts [2] for a linear network.

6. REFERENCES

- [1] Kang, W. N., Kelly, F. P., Lee, N. H. and Williams, R. J. (2007). State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy, preprint.
- [2] Massoulié, L. and Roberts, J. (2000). Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, **15** 185–201.
- [3] Williams, R. J. (1998). An invariance principle for semimartingale reflecting Brownian motions in an orthant. *Queueing Systems*, **30** 5-25.