arXiv:0809.2697v1 [math.PR] 16 Sep 2008

Proportional Fairness and its Relationship with Multi-class Queueing Networks

N. S. Walton*
University of Cambridge

March 15, 2019

Abstract

A network of single server queues with routing is considered. These networks have a product form stationary distribution. A new limit result proves a sequence of such networks converges weakly to a stochastic flow level model. This stochastic model is insensitive. A large deviation principle for the stationary distribution of these queueing networks is found. Its rate function has a dual formulation that coincides with proportional fairness. It is proven that the stationary throughput of the queueing model converges to a proportionally fair allocation.

The queueing models considered have no prescribed optimization structure. Regardless of this, we find a proportionally fair allocation forms an entropy minimizing state of these networks. Proportional fairness occurs as a consequence of a collapse in the state space of the queueing model.

This work combines classical queueing networks with more recent work on stochastic flow level models and proportional fairness. One could view these seemingly different models as the same system described at different levels of granularity: a microscopic, queueing level description; a macroscopic, flow level description and a teleological, optimization description.

1 Introduction

There are various multi-class queueing networks with exponential service requirements that are well understood (12). In this paper we are be interested in viewing a class of such queueing networks at different levels of granularity. To explain what we mean by 'different levels of granularity', let us first outline the example of an electrical network, as described in Doyle and Snell (8).

One could form a Markov chain model of electrons in an electrical network, for example see Kingman (16). At this first level, one explicitly describes the location of particles. Transitions within the network occur rapidly, so perhaps it is more natural to consider results like Ohm's law and Kirchhoff's law which are

*n.s.walton@statslab.cam.ac.uk

concerned with the current flowing through the network. At this second level, one considers the average flow of particles through the network. An electrical network also minimizes energy dissipation, as described by Thomson's principle (primal) and Dirichlet's principle (dual). At this third and final level, one considers the network to be acting as an optimizer.

More generally, we will use the terms microscopic, macroscopic and teleological, as used by Shah and Wischik (25), to classify these different levels of granularity. *Microscopic* refers to a detailed description of a network's state though gives little insight into overall dependence. *Macroscopic* refers to an averaged view of the original network. *Teleological* considers a further abstracted view where the network can be seen to be acting globally as an optimizer.

Here, the aim of viewing networks at these different levels of granularity is to analyse a class of classical queueing networks within the framework of more recent stochastic flow level models (18; 9; 2; 4) and utility optimization formulations of resource allocation (14; 26).

In particular, we will verify a mathematical relationship between multi-class networks of single server queues and proportional fairness. An argument connecting multi-class queueing networks and proportional fairness has been made by Schweitzer (1979) (24), Kelly (1989) (13) and Massoulié and Roberts (1999) (19). This noteworthy argument, presented in section 6, considers the constraints that a queueing network might impose on transfer rates.

We now give an outline of this paper and its key results. In section 2, we introduce notation and the definition of proportional fairness. In section 3, we define our microscopic model. This is a family of multi-class queueing networks where queues have fixed service rates. These networks form a simple mathematical model of document transfer in a packet switching network. We will present known results on the product form stationary distribution of these networks.

In section 4, we describe a convergence technique which we use to form our macroscopic model. This gives the first key result of this paper, theorem 4.1. Theorem 4.1 provides an original limit result of a packet switching model converging to a bandwidth allocation model. We limit from a sequence of multi-class queueing models. Here documents are transferred by discrete packets which are stored and routed through queues. These queueing networks limit to a stochastic flow level model. Here documents are transferred by continuous flows which are dynamically and elastically shared across routes. We call this stochastic flow level model *the spinning network*.

In section 5, we consider properties of our macroscopic model. In particular, we find the spinning network's stationary distribution and, using the results of Bonald and Proutière (4), show that this stationary distribution is insensitive to different document size distributions of the same mean.

In section 6, as mentioned, we present an argument by Schweitzer (24), Kelly (13) and Massoulié and Roberts (19), that relates proportional fairness to multi-class queueing networks. The aim of this section is to develop intuition relevant to following sections.

In section 7, we consider convergence to the teleological description. This

gives the second key result of this paper, theorem 7.1. Theorem 7.1 finds a large deviation principle for the number of documents in transfer across routes for an open multi-class queueing network and also for the spinning network. The rate function found has a primal and dual formulation. The primal form was previously found by Pittel (22). Here the dual formulation is found and coincides with proportionally fair optimization. This dual form of the rate function is considered by Massoulié (17) for two different insensitive stochastic flow level models.

In section 8, we discuss how our proportionally fair, teleological description relates to the behaviour of our original queueing model. Here the third and fourth results of this paper are presented, theorem 8.1 and theorem 8.2. We find the solutions to the primal optimization problem given in theorem 7.1. Using arguments of Pittel (22), theorem 8.1 proves that the state of the multi-class queueing network converges to these solutions and as a result has an optimization characterization. This result found from large deviations is analogous to the heavy traffic, state space collapse results of Bramson (5), Kelly and Williams (15), and Shah and Wischik (25). The fourth theorem, theorem 8.2, proves that as the number of documents in transfer gets larger, the stationary transfer rate of documents converges to a proportionally fair allocation. This result is significant because the queueing networks considered are not designed to optimize but due to this collapse in the state space one sees that these networks asymptotically behave as a proportionally fair optimizer.

Other stochastic models that achieve proportional fairness can be found to inherit properties discussed here. For example, motivated by diffusion approximation behaviour in queueing networks, Kang, Kelly, Lee and Williams (10; 11) considered heavy traffic approximations of stochastic flow level models operating under proportional fairness and found them to have a product form stationary distribution. Also, Massoulié (17) considered a stochastic flow level model operating under balanced fairness, an insensitive bandwidth allocation, and found its stationary distribution to have a large deviations relationship with proportional fairness.

Here we see the advantage in viewing multi-class queueing networks at different levels of granularity is that at each level a different feature of the system can be observed: product form at the microscopic level; insensitivity at the macroscopic level and proportional fairness at the teleological level.

2 Network Structure and Proportional Fairness

Let finite set $\mathcal{J}$ index the set of *queues* in a network. Let $J = |\mathcal{J}|$. A *route* through the network is a non-empty set of queues. Let $\mathcal{I} \subset 2^{\mathcal{J}}$ be the set of routes. Let $I = |\mathcal{I}|$. For each route $i = \{j_1^i, ..., j_{k_i}^i\} \in \mathcal{I}$, we associate an order $(j_1^i, ..., j_{k_i}^i)$. Also we define the set of queue-route incidences, $\mathcal{K} := \{(j, i) : i \in \mathcal{I}, j \in \mathcal{J}, j \in i\}$ and define $K := |\mathcal{K}|$.

Vector $m = (m_{ji} : (j, i) \in \mathcal{K}) \in \mathbb{R}_+^K$ will always be used to refer to either the number or proportion of packets of each route at each queue. Vector $n =$

$(n_i : i \in \mathcal{I}) \in \mathbb{R}_+^I$ will always be used to refer to the number or proportion of documents in transfer across routes.[1] Because each packet in transfer will correspond to a document in transfer we have that,

$$n_i = \sum_{j:j\in i} m_{ji}, \qquad \forall i \in \mathcal{I}.$$

We will always define the number or proportion of packets at a queue to be,

$$m_j := \sum_{i:j\in i} m_{ji}, \qquad \forall j \in \mathcal{J}.$$

For each $n \in \mathbb{Z}_+^I$ we define $S(n) = \{m \in \mathbb{Z}_+^K : \sum_{j:j\in i} m_{ji} = n_i \ \forall i \in \mathcal{I}\}$, the set of queue states achievable given the number of documents in transfer.

With each queue $j \in \mathcal{J}$ we associate a service capacity C_j. A *bandwidth allocation* is a vector $\Lambda(n) = (\Lambda_i(n) : i \in \mathcal{I}) \in \mathbb{R}_+^I$ for each $n \in \mathbb{Z}_+^I$. If in addition $\forall n \in \mathbb{Z}_+^I$,

$$\sum_{i:j\in i} \Lambda_i(n) \le C_j, \qquad \forall j \in \mathcal{J},$$

the bandwidth allocation is said to be *feasible*. A bandwidth allocation corresponds to the rate route i documents are transferred given the number of documents in transfer. Bandwidth allocations form an abstraction of the stationary transfer rate achieved by rate control algorithms used, for example, in the Internet.

A bandwidth allocation $\Lambda^{PF}(n)$ is *proportionally fair* if $\forall n \in \mathbb{R}_+^I$ it solves

$$\begin{aligned}
&\text{maximize} && \sum_{i\in\mathcal{I}} n_i \log \Lambda_i && && (2.1)\\
&\text{subject to} && \sum_{i:j\in i} \Lambda_i \le C_j, && \forall j \in \mathcal{J}, && (2.2)\\
&\text{over} && \Lambda_i \ge 0, && \forall i \in \mathcal{I} && (2.3)
\end{aligned}$$

and if $\Lambda_i^{PF}(n) = 0$ when $n_i = 0$. In the context of communication networks this optimization formulation was first introduced in 1997 (14).

For all vectors $x \in \mathbb{R}^d$ we define $\lfloor x \rfloor := (\lfloor x_1 \rfloor, ..., \lfloor x_d \rfloor)$, the lower integer part of each component. Unless stated otherwise $||x|| = \max_{d'} |x_{d'}|$, the supremum norm. Finally, we will define that for each $m \in \mathbb{Z}_+^K$,

$$\binom{m_j}{m_{ji} \ : i \ni j} = \frac{m_j!}{\prod_{i:j\in i}(m_{ji}!)}.$$

3 A Microscopic Queueing Model

In this section we introduce our microscopic model. This is a multi-class queueing model and forms a simple Markov chain model of document transfer across

[1] The use of m and n as either a number or a proportion will be clear from context.

a packet switching network. Documents wishing to be transferred across a network are broken into a number of packets. These packets are then sent across the network one by one, so that a new packet is sent into the network once the packet in the network has completed its route. We more explicitly define our queueing model in the next three paragraphs. We call it an *open multi-class queueing network with spinning*.

We consider a network of queues indexed by set $\mathcal{J}$. Queues process packets. Each packet moves along a fixed route from the set $\mathcal{I}$. Each queue $j \in \mathcal{J}$ has a fixed service capacity $C_j \in (0, \infty)$ which is then divided between the packets at the queue. Each queue operates under a service discipline that cannot discriminate between the routes used by its packets. So a server may only allocate service between different positions in its queue as a function of the number of packets at the queue. Similarly when joining a queue a packet may only chose its position as a function of the number of packets at that queue.

Documents for transfer along route $i \in \mathcal{I}$ arrive as a Poisson process of rate ν_i. Each route i document consists of a discrete number of packets. This number is independent, finite mean and with distribution equal to random variable X_i. We let $\mu_i = (\mathbb{E}X_i)^{-1}$ and let $\rho_i = \frac{\nu_i}{\mu_i}$ for all $i \in \mathcal{I}$. Each packet has an independent, exponentially distributed mean 1 service requirement at each queue.

Consider route $i \in \mathcal{I}$ with route order $(j_1^i, ..., j_{k_i}^i)$. If we wish to transfer a document across route i, a packet is sent along route i. It will first join queue j_1^i. For $k = 1, ..., k_i - 1$ on departing queue j_k^i a packet will join queue j_{k+1}^i. When a packet leaves its final queue $j(i) = j_{k_i}^i$, a new packet is sent along route i until all packets in the document are transferred.

Equivalently, one could think of each document as a single packet which repeats its route. After completing its route there is some probability it must repeat its route. This probability only depends on the number of repetitions of its route made. This repetition of routes motivates the use of the word 'spinning'.

We now consider the stationary distribution of this model. The following result is a direct consequence of the theorem 3.1 of Kelly (12).

Proposition 3.1. *An open multi-class queueing network with spinning is ergodic iff*

$$\sum_{i:j\in i} \rho_i < C_j, \qquad \forall j \in \mathcal{J}. \tag{3.1}$$

When ergodic, $M = (M_{ji} : (j, i) \in \mathcal{K})$ the number of packets of each route at queues, has stationary distribution,

$$\mathbb{P}(M = m) = B^{-1} \prod_{j\in\mathcal{J}} \left(\binom{m_j}{m_{ji} \ : i \ni j} \prod_{i:j\in i} \left(\frac{\rho_i}{C_j}\right)^{m_{ji}} \right), \tag{3.2}$$

for each $m \in \mathbb{Z}_+^K$*, here we define,*

$$B := \prod_{j \in \mathcal{J}} \left(\frac{C_j}{C_j - \sum_{i \ni j} \rho_i} \right). \tag{3.3}$$

Proof. Allow the state of a packet to be given by the packet's route, the packet's position in its current queue, the total number of repetitions the packet makes of its route and the number of repetitions currently made. From this we have a process of the form described in section 3.1 of (12). Applying theorem 3.1 of (12) to find the stationary distribution and summing over the correct states gains the result. □

The next two corollaries are an immediate consequence of this result.

Corollary 3.1. *$N = (N_i : i \in \mathcal{I})$ the number of documents in transfer has stationary distribution,*

$$\mathbb{P}(N = n) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \rho_i^{n_i}, \qquad \forall\, n \in \mathbb{Z}_+^I, \tag{3.4}$$

where we define,

$$B_n := \sum_{m \in S(n)} \prod_{j \in \mathcal{J}} \left(\binom{m_j}{m_{ji} : i \ni j} \prod_{i : j \in i} \left(\frac{1}{C_j} \right)^{m_{ji}} \right), \qquad \forall\, n \in \mathbb{Z}_+^I. \tag{3.5}$$

Remark 1. *Our definition of an open multi-class queueing network with spinning includes the case where documents consist of a single packet. In this case, as a Poisson process packets enter the network they traverse their route once and then leaves the network. Thus all the results above and all the proportional fairness results in sections 7 and 8 apply to this special case.*

A *closed multi-class queueing network* behaves as an open multi-class queueing network except that document arrivals and departures are forbidden, see (12). So, the network behaves as if there are a fixed number of infinitely large documents in transfer. In the next result we assume the state space of this network is irreducible. [2]

Corollary 3.2. *For a closed multi-class queueing network with $n \in \mathbb{Z}_+^I$ documents in transfer across routes, the number of packets of each route at queues has stationary distribution,*

$$\mathbb{P}_n(M = m) = B_n^{-1} \prod_{j \in \mathcal{J}} \left(\binom{m_j}{m_{ji} : i \ni j} \prod_{i : j \in i} \left(\frac{\rho_i}{C_j} \right)^{m_{ji}} \right), \tag{3.6}$$

for each $m \in \mathbb{Z}_+^K$, where B_n is defined by (3.5).

[2] This condition avoids reducibility issues that can occur from networks in which queues serve a single deterministically chosen packet. In section 4, to simplify the analysis and to satisfy this condition a processor sharing discipline is considered. This assumption is only needed here in corollary 3.2 and in section 4.

4 Limit to a Macroscopic Model

For an electrical network at the macroscopic level, we considered the current through the network and not the explicit behaviour of individual electrons. Similarly for a packet switched network, we may wish to consider the rate documents are transferred through the network and not the explicit position of packets. Stochastic flow level models provide such a model of document transfer without considering packets. In this section we justify how a series of multi-class queueing networks with spinning converges in the Skorohod topology to a stochastic flow level model. Thus we limit from a model where documents are transferred by sending discrete packets to a model where documents are transfered at a dynamic, elastic rate.

A *stochastic flow level model operating under bandwidth allocation* $\Lambda(\cdot)$ is a continuous time Markov chain on $\mathbb{Z}_+^I$ with rates,

$$q(n,n') = \begin{cases} \nu_i & \text{if } n' = n + e_i, \\ \mu_i \Lambda_i(n) & \text{if } n' = n - e_i \text{ and } n_i > 0, \\ 0 & \text{otherwise.} \end{cases} \tag{4.1}$$

e_i is the i-th unit vector in $\mathbb{Z}_+^I$. This model was first considered by Roberts and Massoulié (18).

This model can be interpreted as follows. Documents wishing to be transfered across route i arrive as a Poisson process of rate ν_i. These documents are assumed to have a size that is independent and exponentially distributed with mean μ_i^{-1}. If currently the number of documents in transfer across routes is given by vector $n \in \mathbb{Z}_+^I$ then each document on route i is transferred at rate $\frac{\Lambda_i(n)}{n_i}$. Documents are then processed at this rate until there is a change in the network's state, either by a document transfer being completed and thus leaving the network, or by a document arrival occurring. Thanks to the memoryless property of our process we need not record residual document sizes when an arrival or departure event occurs.

We now introduce the sequence of multi-class queueing networks which we will limit to form our macroscopic model. Consider a sequence of open multi-class queueing networks with spinning, $\{(M^{(c)}(t) : t \in \mathbb{R}_+)\}_{c \in \mathbb{N}}$. These networks have the same routing structure and document arrival processes as described in the last section. We assume for simplicity that each queue is processor sharing, thus $M^{(c)}$ is a Markov chain. We increase the rate packets are transferred through the network. In the c-th network each queue j operates at service rate cC_j. We also increase the size of documents, so as not to increase the rate that documents are transfered through the network. We assume route i documents are geometrically distributed with parameter μ_i/c. We also let $N^{(c)}$ be the stochastic process for the number of documents of each route in transfer, that is, $\forall i \in \mathcal{I}$, $\forall t \in \mathbb{R}_+$,

$$N_i^{(c)}(t) = \sum_{j : j \in i} M_{ji}^{(c)}(t).$$

Let us consider intuitively how these networks limit as $c \to \infty$. Note transitions of packets between queues occur at times of order $O(\frac{1}{c})$. Thus the number of packets sent by a route i document in transfer per unit time is $O(c)$. The probability a packet sent is the final packet is $\frac{\mu_i}{c}$, so the time until this document is transfered is $\frac{\mu_i}{c}O(c) = O(1)$. Thus there is a separation of time scales between document transfer and packet transfer. Between an arrival and departure event the network behaves as a closed queueing network. By the ergodic theorem, for large c, this closed queueing network will behave close to its stationary distribution. Thus between arrival and departure events documents experience a transfer rate determined by the stationary distribution of a closed queueing network. Using the fact that the stationary distribution of $M^{(c)}$ does not depend on c one can determine this transfer rate. When $M^{(c)}$ is stationary, for each route i with final queue j,

$$\lim_{h\to\infty}\frac{1}{h}\mathbb{P}(N^{(c)}(t+h) = n - e_i | N^{(c)}(t) = n)$$
$$= \sum_{\substack{m\in S(n):\\ m_{ji}>0}} \frac{\mu_i}{c}\frac{m_{ji}}{m_j} c C_j \mathbb{P}(M^{(c)} = m | N^{(c)} = n) = \mu_i \mathbb{E}_n\left[\frac{M_{ji}}{M_j} C_j \mathbb{I}[M_j > 0]\right],$$

where $\mathbb{P}_n$ is given by (3.6). Also, the increased rate of packet transfer and the geometric number of packets in a document suggests an exponential distribution limit for the size of documents. Thus is seems plausible a stochastic flow level model would be the limit of these queueing networks and that $\forall n \in \mathbb{Z}_+^I$, the rate of transfer, would be determined by bandwidth allocation,

$$\Lambda_i^{SN}(n) = \mathbb{E}_n\big[\frac{M_{ji}}{M_j} C_j \mathbb{I}[M_j > 0]\big], \qquad \forall i \in \mathcal{I},\ \forall j \in i.^{3}$$

We will show this assertion to be correct. We call the bandwidth allocation found the *spinning allocation* and its associated stochastic flow level model the *spinning network*.

Remark 2. *Here we discuss the spinning allocation as the service rate packets receive at queues in a network of processor sharing queues. One can also check that the spinning allocation is the service rate allocated for any of the queueing disciplines discussed in section 3. For example this includes first come first served and last come first served queueing.*

We can prove that this packet level mechanism converges to a stochastic flow level model. The formal statement of this convergence result is the following theorem.

Theorem 4.1. *For each $c \in \mathbb{N}$, take an open multi-class queueing network with spinning $M^{(c)}$ as described above. Let $N^{(\infty)}$ denote the number of documents in transfer in the spinning network. If*

$$N^{(c)}(0) \Rightarrow N^{(\infty)}(0) \quad as \quad c \to \infty$$

[3] We will show $\Lambda_i^{SN}(n)$ does not depend on j in the next section.

then, in the Skorohod topology[4] *on interval* $[0,1]$,

$$N^{(c)} \Rightarrow N^{(\infty)} \quad as \quad c \to \infty.$$

Due to its technical nature this result is proven in the appendix. The proof uses a coupling argument. The key idea in this proof is to let the internal behaviour of each queueing network between arrival and departure events to be governed by the same process whilst allowing the number of packets sent before a departure to converge almost surely.

Remark 3. *Note that one could use a similar model of a multi-class queueing network with spinning where the number of packets sent is not geometrically distributed and perform this limit. In this way one would model the transfer of arbitrarily distributed documents.*

5 A Macroscopic Stochastic Flow Level Model

In the last section we introduced stochastic flow level models and justified how one such model, the spinning network, formed a macroscopic model of the queueing networks considered in section 3. Now we will discuss some properties of the spinning network. In particular, we show its stationary distribution to be insensitive to different document size distributions.

In the following lemma, we will first prove two facts about the spinning allocation.

Lemma 5.1. *For all* $n \in \mathbb{Z}_+^I$ *and for all* $i \in \mathcal{I}$,

$$\Lambda_i^{SN}(n) = \begin{cases} \frac{B_{n-e_i}}{B_n} & \text{if } n_i > 0 \\ 0 & \text{otherwise,} \end{cases}$$

and where B_n *is defined by (3.5). Moreover,*

$$\sum_{i:j\in i} \Lambda_i^{SN}(n) \le C_j, \quad \forall j \in \mathcal{J},$$

so the spinning allocation is a feasible bandwidth allocation.

Proof. Given stationary distribution (3.6),

$$\begin{aligned}
\Lambda_i^{SN}(n) &= \mathbb{E}_n[\frac{M_{ji}}{M_j} C_j \mathbb{I}[M_j > 0]] \\
&= \frac{1}{B_n} \sum_{\substack{m\in S(n): \\ m_j>0}} \frac{m_{ji}}{m_j} C_j \prod_{l\in\mathcal{J}} \left(\binom{m_l}{m_{lr} : r \ni l} \prod_{r:l\in r} \left(\frac{1}{C_l}\right)^{m_{lr}} \right) \\
&= \frac{1}{B_n} \sum_{m'\in S(n-e_i)} \prod_{l\in\mathcal{J}} \left(\binom{m'_l}{m'_{lr} : r \ni l} \prod_{r:l\in r} \left(\frac{1}{C_l}\right)^{m'_{lr}} \right) = \frac{B_{n-e_i}}{B_n}.
\end{aligned}$$

[4] Here we may define the Skorohod topology with any norm on $\mathbb{R}^I$.

Above we cancelled terms and substituted $m'_{lr} = m_{lr} - 1$ if $(l, r) = (j, i)$ and $m'_{lr} = m_{lr}$ otherwise. Also note,

$$\sum_{i:j\in i} \Lambda_i^{SN}(n) = \sum_{i:j\in i} \mathbb{E}_n[\frac{M_{ji}}{M_j} C_j \mathbb{I}[M_j > 0]] = C_j \mathbb{P}_n(M_j > 0) \leq C_j.$$

□

We can extend the definition of a stochastic flow level model so that the size of incoming documents are independent and of any positive distribution. Information on residual document sizes would be needed for such processes to be Markov. Given this extension a stochastic flow level model with mean document sizes given by $(\mu_i^{-1} : i \in \mathcal{I})$ is *insensitive* if the stationary distribution for the number of documents in transfer is the same as all other stochastic flow level models with the same mean.

One can see that the stationary distribution of an open multi-class queueing network with spinning is unchanged by a change in the mean document size. By increasing the network's service capacity and limiting the discrete document size distribution one can approximate continuous document sizes whilst keeping the stationary distribution unchanged. Thus it is reasonable to think the spinning network may be insensitive.

Bonald and Proutière (3; 4) found that key results on the insensitivity of stochastic flow level models are a consequence of existing results on the insensitivity and reversibility of Whittle networks (27; 23; 20).

Proposition 5.1 (Bonald and Proutière (4))**.** *An ergodic stochastic flow level model operating under bandwidth allocation $\Lambda(\cdot)$ is insensitive iff $\exists$ function $\Phi : \mathbb{Z}^I \to \mathbb{R}_+$ with $\Phi(0) = 1$, $\Phi(n) = 0 \ \forall \ n \notin \mathbb{Z}_+^I$ and*

$$\Lambda_i(n) = \frac{\Phi(n - e_i)}{\Phi(n)}, \qquad \forall n \in \mathbb{Z}_+^I, \quad i \in \mathcal{I},$$

moreover,

$$\pi(n) = \Phi(n) \prod_{i\in\mathcal{I}} \rho_i^{n_i}, \tag{5.1}$$

forms an invariant measure for the number of documents in transfer.

We now find the stationary distribution of the spinning network and show it to be insensitive.

Proposition 5.2. *The spinning network is reversible and insensitive to all document size distributions. The spinning network is ergodic iff*

$$\sum_{i:j\in i} \rho_i < C_j, \quad \forall j \in \mathcal{J}. \tag{5.2}$$

The spinning network has the same stationary distribution as the number of documents in transfer in an open multi-class queueing network with spinning, that is distribution, (3.4).

Proof. Insensitivity and the reversible property are an immediate consequence of proposition 5.1 and lemma 5.1. By proposition 5.1,

$$\pi(n) := B_n \prod_{i \in \mathcal{I}} \rho_i^{n_i}, \quad \forall n \in \mathbb{Z}_+^I,$$

is an invariant measure[5]. The sum of $\pi(\cdot)$ over all states is finite iff the stability condition (5.2) holds. When finite this sum equals B given by (3.3) thus giving stationary distribution (3.4). □

6 Relating Multi-class Queueing Networks to Proportional Fairness

In 1979, Schweitzer (24) studied approximations of closed multi-class queueing networks and considered how asymptotic conditions on such networks might satisfy the Kuhn-Tucker conditions for proportionally fair optimization. In 1989, Kelly (13) studied approximations of closed queueing networks and by analogous analysis considered an optimization formulation similar to proportionally fair optimization. In 1999, Massoulié and Roberts (19) studied a fluid type queueing model and used these same Kuhn-Tucker conditions to deduce proportional fairness. In order to develop intuition and to motivate the following sections we present the argument used in these three papers.

Consider a closed multi-class queueing network as given in section 3. Let n_i be the number of packets on route i; let q_j be the mean sojourn time of a packet at queue j; let $\bar{m}_{ji}$ be the mean number of route i packets at queue j and let Λ_i be the mean sending rate of route i packets into the network.

By Little's law,

$$\Lambda_i q_j = \bar{m}_{ji}, \quad \forall\, (j,i) \in \mathcal{K}. \tag{6.1}$$

Summing over $j \in i$ and rearranging gives,

$$\frac{n_i}{\Lambda_i} - \sum_{j \in i} q_j = 0, \quad \forall i \in \mathcal{I}. \tag{6.2}$$

Since queues are stable we know,

$$\sum_{i: j \in i} \Lambda_i \le C_j, \quad \forall j \in \mathcal{J}. \tag{6.3}$$

One can imagine if (6.3) is not strict then $q_j \approx 0$. Thus approximately,

$$q_j \Big(C_j - \sum_{i: j \in i} \Lambda_i \Big) = 0, \quad \forall j \in \mathcal{J}. \tag{6.4}$$

Also,

$$q_j \ge 0, \quad \forall j \in \mathcal{J} \quad \text{and} \quad \Lambda_i \ge 0, \quad \forall i \in \mathcal{I}. \tag{6.5}$$

[5] For a Markov chain one can verify this by the detailed balance equations.

Interpreting $(q_j : j \in \mathcal{J})$ as Lagrange multipliers, (6.2)-(6.5) are precisely the Kuhn-Tucker conditions for the proportionally fair optimization problem,

$$\max_{\Lambda \geq 0} \quad \sum_{i \in \mathcal{I}} n_i \log \Lambda_i \quad \text{subject to} \quad \sum_{i:j \in i} \Lambda_i \leq C_j \quad \forall j \in \mathcal{J}.$$

So from this one can deduce that $\Lambda_i = \Lambda_i^{PF}(n)\ \forall i \in \mathcal{I}$.

In order to make this argument we needed to assume that the sojourn times of packets did not depend on the route used and that complementary slackness condition (6.4) held. Neither of these conditions need be true in general. In fact we already know that $\Lambda_i = \Lambda_i^{SN}(n)$, the spinning allocation. As the spinning allocation is not Pareto efficient we know that in general $\Lambda_i^{SN}(n) \neq \Lambda_i^{PF}(n)$. Even so it is reasonable to assume $\Lambda_i^{SN}(n) \approx \Lambda_i^{PF}(n)$.

In the following two sections we rigorously prove a relationship between multi-class queueing networks and proportionally fairness. We consider a multi-class network of single server queues, as described in section 5. We let the number of documents in transfer get large but in proportion to some fixed vector $n \in \mathbb{R}_+^I$. We show that these multi-class queueing networks asymptotically allocate service across routes as a proportionally fair optimizer. For example we will show that for all $i \in \mathcal{I}$,

$$\Lambda_i^{SN}(\lfloor hn \rfloor) \xrightarrow[h \to \infty]{} \Lambda_i^{PF}(n).$$

7 Limit to a Teleological Description

In the introduction we discussed how minimizing energy dissipation gave an optimization description of an electrical network. In this section we wish to justify how proportional fairness provides an optimization description for the open multi-class queueing networks discussed in section 3.

To do this we allow the number of documents in transfer to be large and in proportion to some fixed vector. The main result in this section is theorem 7.1 where we prove a large deviation principle for stochastic models with stationary distribution (3.4). This stationary distribution includes the number of documents in transfer for all open multi-class queueing networks discussed in section 3 and the spinning network with any finite mean document size distribution. This approach for queueing networks is similar to that given by Pittel (22), although Pittel does not consider a relationship with proportional fairness in his analysis. In addition, as we consider the large deviations of an insensitive stochastic flow level model the approach is also similar to that taken by Massoulié (17) for balanced fairness.

The large deviation rate function found in theorem 7.1 is,

$$\alpha_\rho(n) = \max_{\Lambda \in \mathbb{R}_+^I} \quad \sum_{i:n_i>0} n_i \log \frac{\Lambda_i}{\rho_i} \quad \text{subject to} \quad \sum_{i:j \in i} \Lambda_i \leq C_j, \ \ \forall j \in \mathcal{J}. \qquad (7.1)$$

As this optimization problem is maximized by the proportionally fair allocation $\Lambda^{PF}(n)$, these queueing models are related to proportionally fair optimization.

To prove theorem 7.1, first we prove a large deviation principle for the stationary distribution (3.2). Stirling's formula finds a rate function $\beta_\rho(\cdot)$. Applying the contraction mapping principle gives the large deviation principle for the number of documents in transfer and finds $\alpha_\rho(\cdot)$ expressed as the primal of a convex optimization problem. We calculate the dual of this optimization problem and find it to be of the form of (7.1).

We start by finding rate function $\beta_\rho(\cdot)$,

Lemma 7.1. *Suppose M is a random variable in $\mathbb{Z}_+^K$ with distribution (3.2). If we take a vector $m \in \mathbb{R}_+^K$ and take $\{d^{(h)}\}_{h\in\mathbb{N}}$ a sequence of vectors in $\mathbb{R}^K$ such that $hm + d^{(h)} \in \mathbb{Z}_+^K$ and $\sup_h |d^{(h)}| < \infty$ then*

$$\lim_{h\to\infty} \frac{1}{h} \log \mathbb{P}(M = hm + d^{(h)}) = -\beta_\rho(m),$$

where we define,

$$\beta_\rho(m) := \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_j>0}} m_{ji} \log \frac{m_{ji} C_j}{m_j \rho_i}. \tag{7.2}$$

Proof. For all $j \in \mathcal{J}$ define, $d_j^{(h)} = \sum_{i:j\in i} d_{ji}^{(h)}$. By Stirling's formula,

$$\begin{aligned}
&\lim_{h\to\infty} \frac{1}{h} \log \mathbb{P}(M = hm + d^{(h)}) \\
&= \lim_{h\to\infty} \frac{1}{h} \Bigg[\sum_{j\in\mathcal{J}} \log (hm_j + d_j^{(h)})! - \sum_{(j,i)\in\mathcal{K}} \log (hm_{ji} + d_{ji}^{(h)})! \\
&\qquad\qquad + \sum_{(j,i)\in\mathcal{K}} (hm_{ji} + d_{ji}^{(h)}) \log \frac{\rho_i}{C_j} \Bigg] \\
&= \lim_{h\to\infty} \frac{1}{h} \Bigg[\sum_{\substack{j\in\mathcal{J}:\\ m_j>0}} \Big((hm_j + d_j^{(h)}) \log(hm_j + d_j^{(h)}) - (hm_j + d_j^{(h)}) \Big) \\
&\qquad - \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_{ji}>0}} \Big((hm_{ji} + d_{ji}^{(h)}) \log(hm_{ji} + d_{ji}^{(h)}) - (hm_{ji} + d_{ji}^{(h)}) \Big) \\
&\qquad\qquad + \sum_{(j,i)\in\mathcal{K}} (hm_{ji} + d_{ji}^{(h)}) \log \frac{\rho_i}{C_j} \Bigg] \\
&= -\lim_{h\to\infty} \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_{ji}>0}} m_{ji} \log \frac{(m_{ji} + \frac{d_{ji}^{(h)}}{h}) C_j}{(m_j + \frac{d_j^{(h)}}{h})\rho_i} = -\beta_\rho(m).
\end{aligned}$$

□

Note that, since $x \log x$ is a continuous function[6], $\beta_\rho(\cdot)$ is a continuous function. Note also lemma 7.1 applies for $d^{(h)} = hm - \lfloor hm \rfloor$. From this lemma the following result is reasonable.

Proposition 7.1. *If M is a random variable in $\mathbb{Z}_+^K$ with distribution (3.2) then, as $h \to \infty$, $\{\frac{M}{h}\}_{h\in\mathbb{N}}$ obeys a large deviation principle with good rate function $\beta_\rho(\cdot)$.*

A proof of this proposition can be found in the appendix. To prove the main theorem of this section we require two lemmas.

Lemma 7.2. *For all $\Lambda \in (0,\infty)^I$,*

$$\inf_{m\in\mathbb{R}_+^K} \beta_\Lambda(m) = \begin{cases} 0 & \text{if} \quad \sum_{i:j\in i} \Lambda_i \le C_j, \ \ \forall j \in \mathcal{J} \\ -\infty & \text{otherwise.} \end{cases}$$

Lemma 7.3. *For all $\Lambda \in (0,\infty)^I$, $\beta_\Lambda(\cdot)$ is a convex function.*

For proofs of lemmas 7.2 and 7.3 see the appendix. We now prove the main theorem of this section. We use the contraction principle, see (7, page 126). Recall that distribution (3.4) is the stationary distribution for the number of documents in transfer for all open multi-class queueing networks discussed in section 3 and the spinning network.

Theorem 7.1. *If N is a random variable in $\mathbb{Z}_+^I$ with distribution (3.4) then as $h \to \infty$, $\{\frac{N}{h}\}_{h\in\mathbb{N}}$ obeys a large deviation principle with good rate function,*

$$\alpha_\rho(n) := \min_{m\in\mathbb{R}_+^K} \sum_{\substack{(j,i)\in\mathcal{K}: \\ m_j>0}} m_{ji} \log \frac{m_{ji}C_j}{m_j\rho_i} \quad \text{subject to} \sum_{j:j\in i} m_{ji} = n_i, \ \forall i \in \mathcal{I} \tag{7.3}$$

$$= \max_{\Lambda\in\mathbb{R}_+^I} \sum_{i\in\mathcal{I}} n_i \log \frac{\Lambda_i}{\rho_i} \quad \text{subject to} \sum_{i:j\in i} \Lambda_i \le C_j, \ \forall j \in \mathcal{J}. \tag{7.4}$$

Proof. Apply the contraction principle to proposition 7.1 using continuous map $f : \mathbb{R}_+^K \to \mathbb{R}_+^I$ such that $f(m) = (\sum_{j:j\in i} m_{ji} : i \in \mathcal{I})$. This gives that $\{\frac{N}{h}\}_{h\in\mathbb{N}}$ obeys a large deviation principle with good rate function,

$$\alpha_\rho(n) = \min_{m\in\mathbb{R}_+^K} \sum_{\substack{(j,i)\in\mathcal{K}: \\ m_j>0}} m_{ji} \log \frac{m_{ji}C_j}{m_j\rho_i} \quad \text{subject to} \quad \sum_{j:j\in i} m_{ji} = n_i, \quad \forall i \in \mathcal{I}.$$

By lemma 7.3, this is a convex optimization problem. Let us calculate its dual

[6] We define $x \log x := 0$ when $x = 0$.

formulation. Taking Lagrange multipliers $\lambda \in \mathbb{R}^I$, its Lagrangian is,

$$\begin{aligned} L(m,\lambda) &= \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_j>0, n_i>0}} m_{ji} \log \frac{m_{ji} C_j}{m_j \rho_i} + \sum_{i:n_i>0} \lambda_i \left(n_i - \sum_{j:j\in i} m_{ji} \right) \\ &= \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_j>0, n_i>0}} m_{ji} \log \frac{m_{ji} C_j}{m_j \rho_i e^{\lambda_i}} + \sum_{i:n_i>0} \lambda_i n_i. \end{aligned}$$

By lemma 7.2,

$$\min_{m\in\mathbb{R}_+^K} L(m,\lambda) = \begin{cases} \sum_{i:n_i>0} n_i \lambda_i & \text{if} \quad \sum_{i:j\in i} \rho_i e^{\lambda_i} \le C_j, \quad \forall j \in \mathcal{J}, \\ -\infty & \text{otherwise,} \end{cases}$$

thus we find dual,

$$\alpha_\rho(n) = \max_{\lambda\in\mathbb{R}^I} \sum_{i:n_i>0} n_i \lambda_i \quad \text{subject to} \quad \sum_{i:j\in i} \rho_i e^{\lambda_i} \le C_j.$$

Substituting $\Lambda_i = \rho_i e^{\lambda_i}$ gives,

$$\alpha_\rho(n) = \max_{\Lambda\in\mathbb{R}_+^I} \sum_{i:n_i>0} n_i \log \frac{\Lambda_i}{\rho_i} \quad \text{subject to} \quad \sum_{i:j\in i} \Lambda_i \le C_j, \quad \forall j \in \mathcal{J}.$$

□

8 Teleological Description

In section 7 we saw that a proportionally fair rate function determined the large deviations behaviour of the stationary distribution of an open multi-class queueing network. In this section we discuss what this means for the behaviour of packets and for the rate of document transfer in these networks. The queueing networks defined in section 3 have no prescribed optimization structure. Even so, as the number of documents gets large, a network will be restricted to its most probable states, because of this the network behaves as an optimizer. This notion of a queueing network collapsing to its most probable states is analogous to the heavy traffic notion of state space collapse (5; 15; 25).

As in section 7, we study the limit as the number of documents in transfer gets large but in proportion to some fixed vector. In this section, we characterize the solutions to the primal problem (7.3). In theorem 8.1, we generalize results from Pittel (22) to show that the state of packets in an open multi-class queueing network with spinning converges in probability to the set of solutions of the primal problem.[7] In corollary 8.1, we show that at each queue the number of packets of each route converges in L^1 to a proportionally fair proportion

[7] Pittel assumes in his analysis that primal problem (7.3) must have a unique solution.

of the number of packets at the queue. Finally in theorem 8.2, we show that the stationary rate documents are transferred through these queueing networks converges to a proportionally fair allocation. We define,

$$\beta(m) := \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_j>0}} m_{ji} \log \frac{m_{ji}C_j}{m_j}, \qquad \forall m \in \mathbb{R}_+^K.$$

We now characterize the solutions to the primal problem (7.3).

Proposition 8.1. *Given $n \in \mathbb{R}_+^I$ and $m^* \in \mathbb{R}_+^K$ such that for all $i \in \mathcal{I}$, $\sum_{j\in i} m^*_{ji} = n_i$ then m^* solves primal problem,*

$$\min_{m\in\mathbb{R}_+^K} \beta(m) \quad \textit{subject to} \quad \sum_{j\in i} m_{ji} = n_i, \quad \forall i \in \mathcal{I} \tag{8.1}$$

iff $\forall (j,i) \in \mathcal{K}$,

$$m^*_{ji}C_j = m^*_j \Lambda_i^{PF}(n). \tag{8.2}$$

Proof. Suppose m^*_{ji} minimizes (8.1), we show (8.2). If $m^*_j = 0$ then $m^*_{ji} = 0$ thus (8.2) holds. Now assume that $m^*_j > 0$, by theorem 7.1,

$$\begin{aligned}
&\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) \\
&= \min_{m\in\mathbb{R}_+^K} \sum_{\substack{(j,i)\in\mathcal{K}:\\ m_j>0, n_i>0}} m_{ji} \log \frac{m_{ji}}{m_j} C_j + \sum_{i:n_i>0} \log \Lambda_i^{PF}(n) \left(n_i - \sum_{j:j\in i} m_{ji} \right),
\end{aligned}$$

and also m^* minimizes this Lagrangian problem. The above expression gives,

$$\min_{m\in\mathbb{R}_+^K} \sum_{j:m_j>0} m_j \sum_{\substack{i:j\in i,\\ n_i>0}} \frac{m_{ji}}{m_j} \log \frac{m_{ji}C_j}{m_j \Lambda_i^{PF}(n)} = 0.$$

Thus by lemma A.5, if $m^*_j > 0$ then $m^*_{ji}C_j = m^*_j \Lambda_i^{SN}(n)$, thus (8.2) holds.

Now we prove the converse. If (8.2) holds then $\forall m^*_j > 0$,

$$\frac{m^*_{ji}C_j}{m^*_j} = \Lambda_i^{PF}(n),$$

therefore,

$$\beta(m^*) = \sum_{i\in\mathcal{I}} \sum_{j:j\in i} m^*_{ji} \log \Lambda_i^{PF}(n) = \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n).$$

□

We define $\forall n \in \mathbb{R}_+^I$,

$$\mathcal{M}(n) = \{m \in \mathbb{R}_+^K : m_{ji} C_j = m_j \Lambda_i^{PF}(n) \; \forall (j,i) \in \mathcal{K}, \; \sum_{j:j\in i} m_{ji} = n_i \; \forall i \in \mathcal{I}\}.$$

In heavy traffic literature $\mathcal{M}(n)$ would be thought of as an *invariant manifold*. Note that if a network of processor sharing queues were in state $m^* \in \mathcal{M}(n)$ then the transfer rate allocated to route i packets would be $\Lambda_i^{PF}(n)$. Since rare events occur in the most likely way, one may expect, as the number of documents in transfer gets large, that the state of the queues in the network will be close to $\mathcal{M}(n)$. Thus the rate of document transfer will be close to a proportionally fair allocation. We use this intuition to prove the next theorem, but first we will require a lemma.

Lemma 8.1. *For all $\epsilon > 0$, $\exists f(\epsilon) > 0$ and $\delta(\epsilon) > 0$ such that $\forall \delta < \delta(\epsilon)$,*

$$\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) + f(\epsilon) \leq \beta^*_{\epsilon,\delta}.$$

Where we define $\forall \epsilon > 0$ and $\delta \geq 0$,

$$\beta^*_{\epsilon,\delta} := \min_{m\in\mathbb{R}_+^K} \beta(m) \quad \textit{subject to} \quad \max_{i\in\mathcal{I}} \Big| \sum_{j\in i} m_{ji} - n_i \Big| \leq \delta$$
$$\textit{and} \quad \inf_{m'\in\mathcal{M}(n)} ||m - m'|| \geq \epsilon.$$

See the appendix for a proof of this result. One could interpret the following result as a state space collapse result (5; 15), as the result says the state space of the queueing network converges to the invariant manifold. Recall all open queueing networks in section 3 have stationary distribution (3.2).

Theorem 8.1. *Let M be any stochastic process on $\mathbb{R}_+^K$ with stationary distribution (3.2) and let $N_i = \sum_{j\in i} M_{ji}$, $\forall i \in \mathcal{I}$ then $\forall n \in \mathbb{R}_+^I$ and $\epsilon > 0$,*

$$\mathbb{P}(\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon \,|\, N = \lfloor hn \rfloor) \xrightarrow[h\to\infty]{} 0.$$

Proof. For all $n \in \mathbb{R}_+^I$ and $\forall \delta > 0$ one has that eventually in h,

$$\begin{aligned}
&\mathbb{P}(\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon \;\Big|\; N = \lfloor hn \rfloor) \qquad (8.3)\\
&= \frac{B}{B_{\lfloor hn \rfloor}} \mathbb{P}(\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon, N = \lfloor hn \rfloor)\\
&\leq \frac{B}{B_{\lfloor hn \rfloor}} \mathbb{P}(\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon, ||\frac{N}{h} - n|| \leq \delta)\\
&\leq \frac{B}{B_{\lfloor hn \rfloor}} \exp\{-h\beta^*_{\epsilon,\delta} + h\frac{f(\epsilon)}{2}\}\\
&\leq \frac{B}{B_{\lfloor hn \rfloor}} \exp\{-h \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) - h\frac{f(\epsilon)}{2}\}.
\end{aligned}$$

We used proposition 7.1 for the second inequality and lemma 8.1 for the final inequality. Selecting any value $m^* \in \mathcal{M}(n)$, one can take sequence $m^{(h)}$ such that $hm^{(h)} \in \mathbb{Z}_+^K$, $\sum_{j:j\in i} m_{ji}^{(h)} = \lfloor hn_i \rfloor \ \forall i \in \mathcal{I}$ and

$$|hm_{ji}^{(h)} - hm_{ji}^*| \leq 2 \quad \forall (j,i) \in \mathcal{K}.$$

So by lemma 7.1 one has,

$$\lim_{h\to\infty} \frac{1}{h} \log \mathbb{P}(M = hm^{(h)}) = \beta(m^*) = -\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n).$$

Thus $\exists h'$ such that $\forall h > h'$,

$$\mathbb{P}(M = hm^{(h)}) \geq \exp\{-h \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) - h\frac{f(\epsilon)}{4}\}. \tag{8.4}$$

Hence one has that, eventually in h,

$$\begin{aligned}
&\mathbb{P}(\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon \big| N = \lfloor hn \rfloor) \\
&\leq \frac{B}{B_{\lfloor hn \rfloor}} \exp\{-h \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) - h\frac{f(\epsilon)}{2}\} \\
&\leq \frac{B}{B_{\lfloor hn \rfloor}} \mathbb{P}(M = hm^{(h)}) e^{-h\frac{f(\epsilon)}{4}} = \mathbb{P}(M = hm^{(h)} | N = \lfloor hn \rfloor) e^{-h\frac{f(\epsilon)}{4}} \xrightarrow[h\to\infty]{} 0.
\end{aligned}$$

Here we used (8.3) for the first inequality and (8.4) for the second inequality. □

We now consider what this notion of result means for the arrangement of packets in queues.

Corollary 8.1. *$\forall n \in \mathbb{R}_+^I$ and $\forall (j,i) \in \mathcal{K}$,*

$$\frac{1}{h}\mathbb{E}\left[\left|M_{ji}C_j - M_j\Lambda_i^{PF}(n)\right| \,\middle|\, N = \lfloor hn \rfloor\right] \xrightarrow[h\to\infty]{} 0.$$

Proof. Take $\epsilon > 0$, by theorem 8.1 and as $m_{ji}C_j = m_j\Lambda_i^{PF}(n)$ $\forall m \in \mathcal{M}(n)$,

$$\begin{aligned}
&\frac{1}{h}\mathbb{E}\left[\left|M_{ji}C_j - M_j\Lambda_i^{PF}(n)\right| \,\middle|\, N = \lfloor hn \rfloor\right] \\
&= \frac{1}{h}\mathbb{E}\left[\left|M_{ji}C_j - M_j\Lambda_i^{PF}(n)\right| \mathbb{I}[\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon] \,\middle|\, N = \lfloor hn \rfloor\right] \\
&+ \frac{1}{h}\mathbb{E}\left[\left|M_{ji}C_j - M_j\Lambda_i^{PF}(n)\right| \mathbb{I}[\inf_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| < \epsilon] \,\middle|\, N = \lfloor hn \rfloor\right] \\
&\leq 2(\sum_{i:j\in i} n_i) C_j \mathbb{P}(\sup_{m'\in\mathcal{M}(n)} ||\frac{M}{h} - m'|| \geq \epsilon \,\big|\, N = \lfloor hn \rfloor) + (1+I)C_j\epsilon \\
&\xrightarrow[h\to\infty]{} 2(1+I)C_j\epsilon.
\end{aligned}$$

Since ϵ is arbitrary the result holds. □

The following result is the formal statement of (6.1) in section 6.

Lemma 8.2. *For all* $(j,i) \in \mathcal{K}$,

$$\Lambda_i^{SN}(n)\mathbb{E}\Big[\ \frac{M_j+1}{C_j}\ \Big|\ N = \lfloor hn \rfloor - e_i\Big] = \mathbb{E}_n M_{ji}. \tag{8.5}$$

Proof. Since $\mathbb{E}_{n-e_i}[\frac{M_j+1}{C_j}]$ is the expected sojourn of a route i packet at queue j the above result is really a statement of Little's law. We can show the result by explicit calculation,

$$\begin{aligned}
\frac{\mathbb{E}_n M_{ji}}{\Lambda_i^{SN}(n)} &= \frac{B_n}{B_{n-e_i}}\frac{1}{B_n}\sum_{\substack{m\in S(n):\\ m_{ji}>0}} m_{ji}\prod_{l\in\mathcal{J}}\left(\begin{pmatrix} m_l \\ m_{lr}\ :r\ni l\end{pmatrix}\prod_{r:l\in r}\left(\frac{1}{C_l}\right)^{m_{lr}}\right)\\
&= \frac{1}{B_{n-e_i}}\sum_{\substack{m\in S(n):\\ m_{ji}>0}} \frac{m_j}{C_j}\prod_{l\in\mathcal{J}}\left(\begin{pmatrix} m_l-\delta_{l,j} \\ m_{lr}-\delta_{lr,ji}\ :r\ni l\end{pmatrix}\prod_{r:l\in r}\left(\frac{1}{C_l}\right)^{m_{lr}-\delta_{lr,ji}}\right)\\
&= \frac{1}{B_{n-e_i}}\sum_{m'\in S(n-e_i)} \frac{m'_j+1}{C_j}\prod_{l\in\mathcal{J}}\left(\begin{pmatrix} m'_l \\ m'_{lr}\ :r\ni l\end{pmatrix}\prod_{r:l\in r}\left(\frac{1}{C_l}\right)^{m'_{lr}}\right)\\
&= \mathbb{E}_{n-e_i}\frac{M_j+1}{C_j}.
\end{aligned}$$

We define $\delta_{x,y} := 1$ for $x = y$ and $\delta_{x,y} := 0$ otherwise, above we cancelled terms and substituted $m'_{lr} = m_{lr} - \delta_{lr,ji}$. □

Consider any of the multi-class queuing networks considered in section 3. Recall that given the number of documents in transfer is n, the stationary rate route i documents are transferred through the network is $\Lambda_i^{SN}(n)$, the spinning allocation. We now prove this rate converges to a proportionally fair bandwidth allocation. In this asymptotic sense these queueing networks transfer documents as a proportionally fair optimizer.

Theorem 8.2. *For all* $n \in \mathbb{R}_+^I$ *and* $\forall i \in \mathcal{I}$,

$$\Lambda_i^{SN}(\lfloor hn \rfloor) \xrightarrow[h\to\infty]{} \Lambda_i^{PF}(n).$$

Proof. If $n_i = 0$ the result is trivially true, so we assume $n_i > 0$. By the arguments in theorem 8.1 and corollary 8.1 one can see that,

$$\frac{1}{h}\mathbb{E}\left[\Big|M_{ji}C_j - M_j\Lambda_i^{PF}(n)\Big|\ \Big|\ N = \lfloor hn \rfloor - e_i\right] \xrightarrow[h\to\infty]{} 0,$$

therefore,

$$\frac{1}{h}\mathbb{E}\left[M_{ji} - \frac{M_j}{C_j}\Lambda_i^{PF}(n)\ \Big|\ N = \lfloor hn \rfloor - e_i\right] \xrightarrow[h\to\infty]{} 0.$$

Summing over $j \in i$ gives,

$$n_i - \frac{\Lambda_i^{PF}(n)}{h} \mathbb{E}\Big[\sum_{j \in i} \frac{M_j}{C_j} \;\Big|\; N = \lfloor hn \rfloor - e_i \Big] \xrightarrow[h \to \infty]{} 0,$$

substituting expression (8.5) gives,

$$n_i - \frac{\Lambda_i^{PF}(n)}{h} \left(\frac{\lfloor n_i h \rfloor}{\Lambda_i^{SN}(\lfloor hn \rfloor)} - \sum_{j \in i} \frac{1}{C_j} \right) \xrightarrow[h \to \infty]{} 0.$$

and finally rearranging gives the result,

$$\Lambda_i^{SN}(\lfloor hn \rfloor) \xrightarrow[h \to \infty]{} \Lambda^{PF}(n).$$

□

9 Conclusion

In this paper we considered a microscopic queueing network as a model of a packet switching network. This network had a product form stationary distribution. We constructed a new limit result to find a stochastic model of document transfer. This macroscopic model had an insensitive stationary distribution. We increased the number of documents in transfer and found proportional fairness to be a teleological description of the queueing network. At each level of granularity, microscopic, macroscopic and teleological, a different feature of the system was respectively found, product form, insensitivity and proportional fairness. The original queueing network had no prescribed notion of optimization; regardless of this we found asymptotically that this network behaved as a proportionally fair optimizer.

Acknowledgments

I am most grateful to Frank Kelly whose insight and enthusiasm has appreciably aided the construction of this paper. I would like to thank Alexandre Proutière for various comments that have greatly aided the content and context of this paper.

A

A.1 Proof of Convergence to the Spinning Network

The proof provided here gives an original proof of a packet level model converging weakly to a stochastic flow level model. Hopefully similar proofs could be developed for other more complex packet level mechanisms.

We require 3 lemmas and a proposition in order to prove theorem 4.1.

Lemma A.1. *If $R(t)$ is a renewal process with mean renewal size γ then, almost surely,*

$$\sup_{t\in[0,\sigma]}\left|\frac{R(ct)}{c}-\frac{t}{\gamma}\right| \xrightarrow[c\to\infty]{} 0, \quad \forall\sigma>0.$$

The proof of this result is an extension of the functional strong law of large number which in turn is a consequence of the strong law of large numbers. See (6, page 106) for a proof of this result.

Lemma A.2. *If X is a positive random variable with $\mathbb{E}X<\infty$ and $(X^{(c)}:c\in\mathbb{N})$ consists of random variables equal in distribution to X then, almost surely,*

$$\lim_{c\to\infty}\frac{X^{(c)}}{c}=0.$$

Proof. $\forall\epsilon\in(0,\infty)\cap\mathbb{Q}$, $\sum_{c\in\mathbb{N}}\mathbb{P}(X^{(c)}>c\epsilon)\leq\epsilon^{-1}\mathbb{E}X+1$ by the Borel-Cantelli lemma the result holds. □

For each $c\in\mathbb{N}$, let $\breve{M}^{(c)}$ be a closed network of processor sharing queues with capacities $C^{(c)}=cC\in\mathbb{R}_+^J$, with $n\in\mathbb{Z}_+^I$ documents in transfer across routes and fixed initial position $m^{(c)}$. Also given $\breve{M}^{(c)}$, let $\breve{R}^{(c)}_{mji}(t)$ be the number of departures of route $i\in\mathcal{I}$ packets from queue $j\in\mathcal{J}$ when the network is in state $m\in S(n)$. For each $n\in\mathbb{Z}_+^I$ let $m_n\in S(n)$ be the state where all packets of every route are in the queue at the beginning of their route. The following proposition will be key to constructing our coupled processes between document arrivals and departures.

Proposition A.1. *There exists a coupling $(\breve{M}^{(c)}:c\in\mathbb{N})$ such that, almost surely, $\forall\sigma>0$, $\forall m\in S(n)$, $(j,i)\in\mathcal{K}$,*

$$\sup_{t\in[0,\sigma]}\left|\breve{R}^{(c)}_{mji}(t)-\frac{m_{ji}}{m_j}C_j\mathbb{I}[m_j>0]\pi_n(m)t\right| \xrightarrow[c\to\infty]{} 0,$$

where $\pi_n(m)$ is the stationary distribution of a closed queueing network, see (3.6). Moreover, the coupled probability space used consists of a countable set of independent uniform $[0,1]$ and exponential mean 1 random variables.

Proof. Take M a closed queueing network with spinning operating under a processor sharing discipline with $n\in\mathbb{Z}_+^I$ packets of each route and capacities $C\in\mathbb{R}_+^J$. Suppose M starts from the fixed initial position $m_n\in S(n)$. Using the jump chain/holding time definition of a Markov chain (21, page 76), we may construct M from a countable set of independent uniform $[0,1]$ and exponential mean 1 random variables. For each $c\in\mathbb{N}$ we define closed queueing network $\hat{M}^{(c)}$ by

$$\hat{M}^{(c)}(t)=M(ct), \qquad \forall t\in[0,1].$$

On the same probability space, we define independently $(\tilde{M}^{(c)}:c\in\mathbb{N})$ to index a sequence of closed queueing networks with $n\mathbb{Z}_+^I$ documents in transfer,

capacities $C^{(c)} = cC$ for each $c \in \mathbb{N}$ and with initial position equal to $(m^{(c)} : c \in \mathbb{N})$. For each $c \in \mathbb{N}$, we can once again form the jumps and holding times of $\tilde{M}^{(c)}$ from one countable set of independent uniform $[0,1]$ random variables and exponential parameter 1 random variables. Note both $\breve{M}^{(c)}$ and $\hat{M}^{(c)}$ are both positive recurrent, finite state space Markov chains with the same generator.

Let $\breve{U}^{(c)}$ be the first time $\tilde{M}^{(c)}$ hits state m_n. We define our coupled process to be,

$$\breve{M}^{(c)}(t) := \begin{cases} \hat{M}^{(c)}(t - \breve{U}^{(c)}) & \text{if} \quad t \geq \breve{U}^{(c)}, \\ \tilde{M}^{(c)}(t) & \text{if} \quad t < \breve{U}^{(c)}. \end{cases}$$

Note by the positive recurrence of $\tilde{M}^{(c)}$ and because given the initial position of $\tilde{M}^{(1)}$ equals $m^{(c)}$, $c\breve{U}^{(c)}$ is equal in distribution to $\breve{U}^{(1)}$ one has that,

$$\mathbb{E}[c\breve{U}^{(c)}] = \mathbb{E}[\breve{U}^{(1)} | \tilde{M}^{(1)}(0) = m^{(c)}] \leq \max_{m \in S(n)} \mathbb{E}[\breve{U}^{(1)} | \tilde{M}^{(1)}(0) = m] < \infty.$$

Thus by lemma A.2, almost surely,

$$\breve{U}^{(c)} = \sum_{m \in S(n)} \breve{U}^{(c)} \mathbb{I}[\tilde{M}^{(c)}(0) = m] \xrightarrow[c \to \infty]{} 0. \tag{A.1}$$

Also, conditional on $\tilde{M}^{(1)} = m^{(c)}$, $\forall m \in S(n)$ and $(j,i) \in \mathcal{K}$, $\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})$ is equal in distribution to $\breve{R}^{(1)}_{mji}(\breve{U}^{(1)})$. By positive recurrence,

$$\mathbb{E}\breve{R}^{(c)}_{mji}(\breve{U}^{(c)}) \leq \max_{m \in S(n)} \mathbb{E}[\breve{R}^{(1)}_{mji}(\breve{U}^{(1)}) | \tilde{M}^{(1)}(0) = m] < \infty.$$

Once again by lemma A.2, almost surely,

$$\frac{\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})}{c} \xrightarrow[c \to \infty]{} 0. \tag{A.2}$$

For closed queueing network M let $R_{mji}(t)$ denote the number of route $i \in \mathcal{I}$ packets to leave queue $j \in \mathcal{J}$ when the network is in state $m \in S(n)$ by time t. Note R_{mji} is a renewal process as for any Markov chain the process that gives the current state of the chain and the state after the next jump is also a Markov chain. For this second Markov chain the stationary probability of the current state being m and the next state being the result of a route i packet leaving queue j is,

$$\frac{C_j}{\sum_{l:m_l>0} C_l} \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0] \pi_n(m).$$

Thus by Little's law renewal process R_{mji} has drift,

$$C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0] \pi_n(m).$$

Also,

$$R_{mji}(ct) = \breve{R}^{(c)}_{mji}(t + \breve{U}^{(c)}) - \breve{R}^{(c)}_{mji}(\breve{U}^{(c)}). \tag{A.3}$$

Thus $\breve{R}^{(c)}_{mji}$ will be governed by the renewal theorem for R_{mji} since $\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})$ and $\breve{U}^{(c)}$ are small. So we see, almost surely, for all $\sigma > 0$,

$$\lim_{c\to\infty} \sup_{t\in[0,\sigma]} \left| \frac{\breve{R}^{(c)}_{mji}(t)}{c} - C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0]\pi_n(m)t \right|$$

$$\leq \lim_{c\to\infty} \left[\frac{\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})}{c} + C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0]\pi_n(m)\breve{U}^{(c)} \right.$$

$$\left. + \sup_{t\in[\breve{U}^{(c)}\wedge\sigma,\sigma]} \left| \frac{R_{mji}(c(t-\breve{U}^{(c)}))}{c} + \frac{\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})}{c} - C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0]\pi_n(m)t \right| \right]$$

$$\leq \lim_{c\to\infty} \left[0 + 0 + \frac{\breve{R}^{(c)}_{mji}(\breve{U}^{(c)})}{c} + C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0]\pi_n(m)\breve{U}^{(c)} \right.$$

$$\left. + \sup_{t\in[0,\sigma]} \left| \frac{R_{mji}(ct)}{c} - C_j \frac{m_{ji}}{m_j} \mathbb{I}[m_j > 0]\pi_n(m)t \right| \right]$$

$$= 0.$$

Here we substituted (A.3). Also we used (A.1) and (A.2) for the second inequality and last equality, and used the lemma A.1 for the last equality. This gives the required result. □

Lemma A.3. *For $c \in \mathbb{N}$, take $D^{(c)} : \mathbb{R}_+ \to \mathbb{R}_+$ an increasing function and $D^{(\infty)}(t) = \Lambda t$ for $\Lambda > 0$. Take sequence $\{y^{(c)}\}_{c\in\mathbb{N}\cup\{\infty\}}$. If $D^{(c)}(\cdot)$ converges uniformly on compacts to $D^{(\infty)}(\cdot)$ and $y^{(c)} \to y^{(\infty)}$ as $c \to \infty$ then*

$$\inf\{t \geq 0 : D^{(c)}(t) = y^{(c)}\} \xrightarrow[c\to\infty]{} \inf\{t \geq 0 : D^{(\infty)}(t) = y^{(\infty)}\} = \frac{y^{(\infty)}}{\Lambda}.$$

Proof. Take $\epsilon > 0$ and $\sigma > \frac{y^{(\infty)}+2\epsilon}{\Lambda}$ and chose c' such that $\forall c \geq c'$,

$$\sup_{t\in[0,\sigma]} |D^{(c)}(t) - D^{(\infty)}(t)| \leq \epsilon, \quad |y^{(c)} - y^{(\infty)}| < \epsilon,$$

then we have that,

$$D^{(c)}(\frac{y^{(\infty)}}{\Lambda} - \frac{2\epsilon}{\Lambda}) \leq D^{(\infty)}(\frac{y^{(\infty)}}{\Lambda} - \frac{2\epsilon}{\Lambda}) + \epsilon \leq y^{(\infty)} - \epsilon < y^{(c)},$$

therefore,

$$\inf\{t \geq 0 : D^{(c)}(t) = y^{(c)}\} > \frac{y^{(\infty)}}{\Lambda} - \frac{2\epsilon}{\Lambda}.$$

Similarly one can see that,

$$\inf\{t \geq 0 : D^{(c)}(t) = y^{(c)}\} < \frac{y^{(\infty)}}{\Lambda} + \frac{2\epsilon}{\Lambda}.$$

This gives the result. □

Proof of theorem 4.1. We prove the result via a coupling argument. We will prove by induction that a coupled processes $\{\bar{N}^{(c)}\}_{c\in\mathbb{N}\cup\{\infty\}}$ eventually make the same jumps on interval $[0,1]$ and that their associated jump times converge almost surely. We prove that the induction hypothesis contained in the next paragraph is satisfied on a single probability space that consists of a countable collection of independent uniform $[0,1]$ and independent exponential parameter 1 random variables.

For $k \in \mathbb{N}$, the induction hypothesis is: there exists a sequence of open queueing networks with spinning, $\bar{M}^{(c)}\ \forall c \in \mathbb{N}$, defined on $[0,\tau^{k,(c)}]$ and spinning network $N^{(\infty)}$ defined on $[0,\tau^{k,(\infty)}]$. Where for $c \in \mathbb{N}\cup\{\infty\}$ and $\kappa \le k$, $\tau^{\kappa,(c)}$ gives the κ-th document arrival or departure event of the c-th network. We also define for $\tau^{\kappa,(c)} = 0$ for $\kappa = 0$. Let process $\bar{N}^{(c)}$ be the number of documents of each route in the c-th network. We have that, almost surely, there exists a random variable c_k such that $\forall c \ge c_k$ and $\forall \kappa \le k$,

$$\bar{N}^{(c)}(\tau^{\kappa,(c)}) = \bar{N}^{(\infty)}(\tau^{\kappa,(\infty)}), \tag{A.4}$$

and $\forall \kappa \le k$,

$$\tau^{\kappa,(c)} \xrightarrow[c\to\infty]{} \tau^{\kappa,(\infty)}. \tag{A.5}$$

This completes the inductive hypothesis. The hypothesis says eventually the first k jumps are identical and the associated jump times converge.

Take a probability space consisting of a countable independent set of uniform $[0,1]$ random variables and exponential parameter 1 random variables. We prove we can construct a process that satisfies the induction hypothesis on this probability space.

We can prove the hypothesis is true for $k = 0$. By the Skorohod Representation theorem with a single uniform $[0,1]$ random variable, we may chose $(\bar{N}^{(c)}(0) : c \in \mathbb{N}\cup\{\infty\})$ such that, almost surely,

$$\bar{N}^{(c)}(0) \xrightarrow[c\to\infty]{} \bar{N}^{(\infty)}(0).$$

Take $\forall c \in \mathbb{N}$, $f^{(c)} : \mathbb{Z}_+^I \times [0,1] \to \mathbb{Z}_+^K$ such that,

$$\mathbb{P}(f^{(c)}(n,U) = m) = \mathbb{P}(M^{(c)} = m | N^{(c)} = n), \quad \forall m \in \mathbb{Z}_+^K,\ n \in \mathbb{Z}_+^I,$$

where U is a uniform [0,1] random variable. From $\bar{N}^{(c)}(0)$ and an additional independent uniform random variable U we define,

$$\bar{M}^{(c)}(0) = f^{(c)}(\bar{N}^{(c)}(0), U), \quad \forall c \in \mathbb{N}.$$

$\bar{M}^{(c)}(0)$ has with distribution equal to $M^{(c)}(0)$, (A.4) holds and as $\tau^{0,(c)} := 0$ (A.5) holds. So the inductive hypothesis is true for $k = 0$.

Suppose the inductive hypothesis is true for $k-1$ for some $k > 0$. We will prove the inductive hypothesis holds for k. Take the coupling for $k-1$, $\bar{M}^{(c)}$ for $c \ge c_{k-1}$ and $\bar{N}^{(\infty)}$, from the induction hypothesis. For each $i \in \mathcal{I}$ take on the same probability space $(Y_i^{k,(c)} : c \in \mathbb{N}\cup\{\infty\})$ to be random variables mutually

independent over $i \in \mathcal{I}$ with, $Y_i^{k,(c)}$ geometrically distributed with parameter $\frac{\mu_i}{c}$ for each $i \in \mathcal{I}$ and $c \in \mathbb{N}$, $Y_i^{k,(\infty)}$ exponentially distributed with parameter μ_i for each $i \in \mathcal{I}$ and such that, almost surely,

$$\frac{Y_i^{k,(c)}}{c} \xrightarrow[c\to\infty]{} Y_i^{k,(\infty)}, \qquad \forall i \in \mathcal{I}.$$

By the Skorohod representation theorem forming this requires I independent uniform random variables.

We set $n^{k-1} = \bar{N}^{(\infty)}(\tau^{k-1,(\infty)})$. Take from our probability space a countable set of uniform and exponential random variables to construct,

$$(\breve{M}^{k,(c)} : c \in \{c_{k-1}, c_{k-1}+1, ..., \infty\}),$$

as taken from proposition A.1 with initial position,

$$(\breve{M}^{k,(c)}(0) : c \in \mathbb{N}) := (\bar{M}^{(c)}(\tau^{k-1,(c)}) : c \in \mathbb{N}).$$

For $\breve{M}^{k,(c)}$, let $D_i^{k,(c)}(t)$ be the number departures of route i packets by time t from queue $j(i)$, the final queue on route $i \in \mathcal{I}$. Let $D_i^{k,(\infty)}(t) = \Lambda_i^{SN}(n^{k-1})t$. Define $\forall c \geq c_{k-1}$ and $c = \infty$,

$$\zeta_i^{k,(c)} = \inf\{t \geq 0 : \ D_i^{k,(c)}(t) = Y_i^{k,(c)}\}.$$

For $c \geq c_{k-1}$ and $c = \infty$, $\zeta_i^{k,(c)}$ would be the first time a route $i \in \mathcal{I}$ document leaves the network if all other document arrivals and departures are not permitted. Define $\forall i \in \mathcal{I}$, η_i^k to be an independent exponentially distributed random variable with parameter ν_i. Define,

$$\tau^{k,(c)} = \tau^{k-1,(c)} + (\min_{i\in\mathcal{I}} \zeta_i^{k,(c)}) \wedge (\min_{i\in\mathcal{I}} \eta_i^k), \quad \forall c \geq c_{k-1} \text{ and } c = \infty.$$

$\tau^{k,(c)}$ gives the time of the k-th document arrival or departure event and the term $\zeta_i^{k,(c)}$ or η_i^k that achieves the minimum determines whether a document arrival or departure occurs.

So, for $c \geq c_{k-1}$, we extend $\bar{M}^{(c)}$ to include $t \in [\tau^{k-1,(c)}, \tau^{k,(c)}]$ as follows,

$$\bar{M}^{(c)}(t) = \begin{cases} \breve{M}^{k,(c)}(t - \tau^{k-1,(c)}) & \text{if } t \in [\tau^{k-1,(c)}, \tau^{k,(c)}), \\ \breve{M}^{k,(c)}(\{\tau^{k,(c)} - \tau^{k-1,(c)}\}^-) - e_{j(i)i} & \text{if } t = \tau^{k,(c)} \\ & \text{and } \tau^{k,(c)} - \tau^{k-1,(c)} = \zeta_i^{k,(c)}, \\ \breve{M}^{k,(c)}(\{\tau^{k,(c)} - \tau^{k-1,(c)}\}^-) + e_{j_1^i i} & \text{if } t = \tau^{k,(c)} \\ & \text{and } \tau^{k,(c)} - \tau^{k-1,(c)} = \eta_i^k. \end{cases} \tag{A.6}$$

Here we define $M(t^-) := \lim_{s\searrow t} M(s)$. Thus for $c \geq c_{k-1}$ and $c = \infty$,

$$\bar{N}^{(c)}(t) = \begin{cases} n^{k-1} & \text{if } t \in [\tau^{k-1,(c)}, \tau^{k,(c)}), \\ n^{k-1} - e_i & \text{if } t = \tau^{k,(c)} \text{ and } \tau^{k,(c)} - \tau^{k-1,(c)} = \zeta_i^{k,(c)}, \\ n^{k-1} + e_i & \text{if } t = \tau^{k,(c)} \text{ and } \tau^{k,(c)} - \tau^{k-1,(c)} = \eta_i^k. \end{cases} \tag{A.7}$$

Here we have constructed, on the same probability space, the spinning network and a multi-class queueing network with spinning, up to the k-th jump.

We now prove these networks have the convergence properties given by the induction hypothesis. Note for each $i \in \mathcal{I}$ that $D_i^{k,(c)}(\cdot)$ is the sum of renewal processes,

$$D_i^{k,(c)}(t) = \sum_{m \in S(n^{k-1})} \bar{R}_{mj(i)i}^{k,(c)}(t)$$

where $\bar{R}_{mji}^{k,(c)}(t)$ gives, as in proposition A.1, the number of times a route $i \in \mathcal{I}$ packet departs queue $j \in \mathcal{J}$ when the system is in state $m \in S(n^{k-1})$ by time t. Thus by proposition A.1, almost surely, $\forall \sigma > 0$, $\forall i \in \mathcal{I}$,

$$\begin{aligned}
&\sup_{t\in[0,\sigma]} \left| \frac{D_i^{k,(c)}(t)}{c} - \Lambda_i^{SN}(n^{k-1})t \right| \\
&\leq \sum_{m \in S(n^{k-1})} \sup_{t\in[0,\sigma]} \left| \frac{\bar{R}_{mj(i)i}^{k,(c)}(t)}{c} - \frac{m_{j(i)i}}{m_{j(i)}} C_{j(i)} \pi_{n^{k-1}}(m) \mathbb{I}[m_{j(i)} > 0] t \right| \xrightarrow[c\to\infty]{} 0.
\end{aligned}$$

Thus by lemma A.3, almost surely, $\forall i \in \mathcal{I}$,

$$\begin{aligned}
\zeta_i^{k,(c)} &= \inf\{t \geq 0 : \frac{D_i^{k,(c)}(t)}{c} = \frac{Y_i^{k,(c)}}{c}\} \\
&\xrightarrow[c\to\infty]{} \inf\{t \geq 0 : \Lambda_i^{SN}(n^{k-1})t = Y_i^{k,(\infty)}\} = \zeta_i^{k,(\infty)}.
\end{aligned} \tag{A.8}$$

Almost surely, the independent exponential terms $\{\zeta_i^{k,(\infty)}, \eta_i^k : i \in \mathcal{I}\}$ are distinct. Thus there is a random variable $c_k > c_{k-1}$ such that $\forall c \geq c_k$ the term achieving the minimum in the set $\{\zeta_i^{k,(c)}, \eta_i^k : i \in \mathcal{I}\}$ corresponds to the term achieving the minimum in the set $\{\zeta_i^{k,(\infty)}, \eta_i^k : i \in \mathcal{I}\}$. So, by this and expression (A.8) almost surely, $\forall c > c_k$,

$$\bar{N}^{(c)}(\tau^{k,(c)}) = \bar{N}^{(\infty)}(\tau^{k,(\infty)}), \tag{A.9}$$

$$\text{and} \qquad \tau^{k,(c)} \xrightarrow[c\to\infty]{} \tau^{k,(\infty)}. \tag{A.10}$$

Since for this, and thus for every, induction step we require a countable number of independent random variables, the induction hypothesis holds $\forall k \in \mathbb{N}$ on the proposed probability space.

Since the rates of all the queueing networks considered are bounded we know that, almost surely,

$$\tau_k^{(c)} \xrightarrow[k\to\infty]{} \infty, \quad \forall c \in \mathbb{N} \cup \{\infty\}.$$

By this and our induction argument the processes $\{\bar{M}^{(c)}\}_{c\in\mathbb{N}}$ and $\bar{N}^{(\infty)}$ are defined for all times in $\mathbb{R}_+$. For all $c \in \mathbb{N}$ these are multi-class queueing networks with spinning and the spinning network for $c = \infty$.

We now work to prove almost sure convergence in the Skorohod topology. For each event we may define the random variable,

$$\Gamma := \max\{k : \tau^{k,(c)} < 1\}.$$

Since (A.9) holds, almost surely we can find c'_Γ such that $\forall c \geq c'_\Gamma$,

$$N^{k,(c)}(\tau^{k,(c)}) = N^{k,(\infty)}(\tau^{k,(\infty)}), \quad \tau^{k,(c)} \in (0,1) \qquad \forall k \leq \Gamma,$$
$$\text{and} \qquad \tau^{\Gamma+1,(c)} \in (1,\infty).$$

For all $c \in \mathbb{N}$ take map $\lambda^{(c)} : [0,1] \to [0,1]$ with $\lambda^{(c)}(0) := 0$, $\lambda^{(c)}(1) := 1$,

$$\lambda^{(c)}(\tau^{k,(c)}) := \tau^{k,(\infty)}, \quad \forall k \leq \Gamma,$$

else define $\lambda^{(c)}(t)$ by linearly interpolating between these terms. So $\forall c \geq c'_\Gamma$,

$$\sup_{t\in[0,1]} |\lambda^{(c)}(t) - t| = \sup_{k\in\{1,..,\Gamma\}} |\tau_k^{(c)} - \tau_k^{(\infty)}|,$$
$$\text{and} \qquad \sup_{t\in[0,1]} ||N^{(\infty)}(\lambda^{(c)}(t)) - N^{(c)}(t)|| = 0,$$

where $||\cdot||$ is any norm on $\mathbb{R}^I$. Take $d(\cdot,\cdot)$ to be the Skorohod metric on $D^I[0,1]$, see (1, page 124). By (A.9) and (A.10) almost surely, $\forall c \geq c_\Gamma$,

$$d(N^{(c)}, N^{(\infty)}) \leq \sup_{t\in[0,1]} |\lambda^{(c)}(t) - t| \vee \sup_{t\in[0,1]} ||N^{(\infty)}(\lambda^{(c)}(t)) - N^{(c)}(t)||$$
$$= \sup_{k\in\{1,..,\Gamma\}} |\tau_k^{(c)} - \tau_k^{(\infty)}| \xrightarrow[c\to\infty]{} 0.$$

Thus since in this probability space these coupled processes converge almost surely in the Skorohod topology, the processes must converge weakly in the Skorohod topology. □

A.2 Proof of Proposition 7.1

To prove proposition 7.1 we need to extend the lemma 7.1 to make a large deviation principle. The following lemmas will be useful for this. The proofs of both lemmas are simple, calculus arguments.

Lemma A.4. *For all $m \in \mathbb{R}_+^K$ and $\forall j \in \mathcal{J}$ with $m_j > 0$, we have*

$$\max_{\substack{\theta_i>0:\\ j\in i}} \{\sum_{i:j\in i} \theta_i m_{ji} : \sum_{i:j\in i} \rho_i e^{\theta_i} = C_j\} = \sum_{i:j\in i} m_{ji} \log \frac{m_{ji} C_j}{m_j \rho_i}.$$

Lemma A.5. *For all $\Lambda \in (0,\infty)^I$ and for each $j \in \mathcal{J}$, we have*

$$\min_{\substack{p_i>0:\\ j\in i}} \{\sum_{i:j\in i} p_i \log \frac{p_i C_j}{\Lambda_i} : \sum_{i:j\in i} p_i = 1\} = \log \frac{C_j}{\sum_{i:j\in i} \Lambda_i},$$

and the minimum is attained by $p_i = \frac{\Lambda_i}{\sum_{r:j\in r} \Lambda_r}$.

We now prove proposition 7.1. The result makes use of the upper bound of the Gärtner-Ellis theorem, see (7, Page 44).

Proof of Proposition 7.1. First we prove the lower bound. Take $O \subset \mathbb{R}_+^K$ open, for all $m \in O \; \exists \; h'$ such that $\forall \; h > h'$, $\frac{\lfloor hm \rfloor}{h} \in O$. By this and lemma 7.1 we have that,

$$\liminf_{h\to\infty} \frac{1}{h} \log \mathbb{P}(\frac{M}{h} \in O) \geq \sup_{m\in O} \liminf_{h\to\infty} \frac{1}{h} \log \mathbb{P}(M = \lfloor hm \rfloor) = - \inf_{m\in O} \beta_\rho(m).$$

This gives the required lower bound.

Now we prove the upper bound. M has moment generating function,

$$\mathbb{E}e^{\sum_{(j,i)\in\mathcal{K}} \theta_{ji} M_{ji}} = \begin{cases} \prod_{j\in\mathcal{J}} \left(\frac{C_j - \sum_{i:j\in i} \rho_i}{C_j - \sum_{i:j\in i} \rho_i e^{\theta_{ji}}} \right) & \text{if } \sum_{i:j\in i} \rho_i e^{\theta_{ji}} < C_j, \; \forall j \in \mathcal{J}, \\ \infty & \text{otherwise.} \end{cases}$$

Thus,

$$F(\theta) := \lim_{h\to\infty} \frac{1}{h} \log \mathbb{E}e^{\sum_{(j,i)\in\mathcal{K}} h\theta_{ji} \frac{M_{ji}}{h}} = \begin{cases} 0 & \text{if } \sum_{i:j\in i} \rho_i e^{\theta_{ji}} < C_j, \; \forall j \in \mathcal{J}, \\ \infty & \text{otherwise.} \end{cases}$$

Let $F^*(\cdot)$ be the Legendre-Fenchel transform of $F(\cdot)$. By lemma A.4, $\forall m \in \mathbb{R}_+^K$,

$$\begin{aligned} F^*(m) &= \sup_{\theta\in\mathbb{R}^K} \{ \sum_{(j,i)\in\mathcal{K}} \theta_{ji} m_{ji} : \sum_{i:j\in i} \rho_i e^{\theta_{ji}} < C_j, \; \forall j \in \mathcal{J} \} \\ &= \sum_{j\in\mathcal{J}} \sup_{\substack{\phi_i > 0: \\ j\in i}} \{ \sum_{i:j\in i} \phi_i m_{ji} : \sum_{i:j\in i} \rho_i e^{\phi_i} < C_j \} \\ &= \sum_{j:m_j>0} \sum_{i:j\in i} m_{ji} \log \frac{m_{ji} C_j}{m_j \rho_i} = \beta_\rho(m). \end{aligned}$$

Thus the Gärtner-Ellis theorem gives for all closed sets $C \subset \mathbb{R}_+^K$,

$$\limsup_{h\to\infty} \frac{1}{h} \log \mathbb{P}(\frac{M}{h} \in C) \leq - \inf_{m\in C} \beta_\rho(m).$$

Finally, we prove $\beta_\rho(\cdot)$ is a good rate function. $\beta_\rho(\cdot)$ is continuous with values in $\mathbb{R}_+$ and so is a rate function. By lemma A.5, $\forall \alpha \geq 0$, $\forall \; m \in \mathbb{R}_+^K$ if

$$m_j > \frac{\alpha}{\log C_j - \log \sum_{i:\; j\in i} \rho_i},$$

then

$$\beta_\rho(m) = \sum_{j:m_j>0} m_j \sum_{i:j\in i} \frac{m_{ji}}{m_j} \log \frac{m_{ji} C_j}{m_j \rho_i} \geq \sum_{j:m_j>0} m_j \log \frac{C_j}{\sum_{i:j\in i} \rho_i} > \alpha.$$

Thus,

$$\{m \in \mathbb{R}_+^K : \; \beta_\rho(m) \leq \alpha\} \subset \{m \in \mathbb{R}_+^K : 0 \leq m_j \leq \frac{\alpha}{\log C_j - \log \sum_{i:j\in i} \rho_i}\}.$$

So all level sets are compact and hence $\beta_\rho(\cdot)$ is a good rate function. □

A.3 Proofs of Additional Lemmas

Proof of Lemma 7.2. By lemma A.5,

$$\begin{aligned}\inf_{m\in\mathbb{R}_+^K}\beta_\Lambda(m) &= \inf_{m\in\mathbb{R}_+^K}\sum_{j:m_j>0} m_j \sum_{i:j\in i}\frac{m_{ji}}{m_j}\log\frac{m_{ji}C_j}{m_j\Lambda_i}\\ &= \inf_{m'\in\mathbb{R}_+^J}\sum_{j:m'_j>0} m'_j\log\frac{C_j}{\sum_{i:j\in i}\Lambda_i} = \begin{cases}0 & \text{if } \sum_{i:j\in i}\Lambda_i \le C_j,\\ -\infty & \text{otherwise.}\end{cases}\end{aligned}$$

□

Proof of Lemma 7.3. We first show $\beta_\Lambda(\cdot)$ is convex on $(0,\infty)^K$. One can see $\forall m\in(0,\infty)^K$ and $\forall (j,i),\ (l,r)\in\mathcal{K}$ that,

$$\frac{\partial^2\beta_\Lambda}{\partial m_{ji}\partial m_{lr}}(m) = \frac{1}{m_{ji}}\mathbb{I}[(j,i)=(l,r)] - \frac{1}{m_j}\mathbb{I}[l=r].$$

So for all vectors $a\in\mathbb{R}^K\backslash\{0\}$ the second partial derivative in direction a is,

$$\frac{\partial^2\beta_\Lambda}{\partial a^2} = \sum_{\substack{(j,i)\in\mathcal{K},\\(l,r)\in\mathcal{K}}} a_{ji}a_{lr}\frac{\partial^2\beta_\Lambda}{\partial m_{ji}\partial m_{lr}}(m) = \sum_{(j,i)\in\mathcal{K}}\frac{a_{ji}^2}{m_{ji}} - \sum_{j\in\mathcal{J}}\sum_{i:j\in i}\sum_{r:j\in r}\frac{a_{ji}a_{jr}}{m_j}.$$

We will prove that this expression is positive. Note that for fixed $j\in\mathcal{J}$,

$$\sum_{i:j\in i} a_{ji}^2\frac{m_j}{m_{ji}} = \sum_{i:j\in i}\sum_{r:j\in r}\frac{1}{2}\left(a_{ji}^2\frac{m_{jr}}{m_{ji}} + a_{jr}^2\frac{m_{ji}}{m_{jr}}\right).$$

For all $i\neq r$, $a_{ji}^2x + a_{jr}^2\frac{1}{x}$ is convex on $(0,\infty)$ and has its minimum at $x=|\frac{a_{jr}}{a_{ji}}|$. So,

$$\sum_{i:j\in i} a_{ji}^2\frac{m_j}{m_{ji}} \ge \sum_{i:j\in i}\sum_{r:j\in r}|a_{ji}||a_{jr}| \ge \sum_{i:j\in i}\sum_{r:j\in r}a_{ji}a_{jr} \ge 0.$$

Dividing by m_j and summing over $j\in\mathcal{J}$ gives that,

$$\frac{\partial^2\beta_\Lambda}{\partial a^2} = \sum_{j\in\mathcal{J}}\left(\sum_{i:j\in i}\frac{a_{ji}^2}{m_{ji}} - \sum_{i:j\in i}\sum_{r:j\in r}\frac{a_{ji}a_{jr}}{m_j}\right) \ge 0$$

This proves $\beta_\Lambda(\cdot)$ is convex on $(0,\infty)^K$.

Take $m,\bar{m}\in\mathbb{R}_+^K$ and two sequences $m^{(h)},\bar{m}^{(h)}\in(0,\infty)^K$ that converge to m and $\bar{m}$ respectively. By continuity of $\beta_\Lambda(\cdot)$, $\forall\theta\in(0,1)$,

$$\begin{aligned}\beta_\Lambda(\theta m+(1-\theta)\bar{m}) &= \lim_{h\to\infty}\beta_\Lambda(\theta m^{(h)}+(1-\theta)\bar{m}^{(h)})\\ &\le \lim_{h\to\infty}\theta\beta_\Lambda(m^{(h)})+(1-\theta)\beta_\Lambda(\bar{m}^{(h)})\\ &= \theta\beta_\Lambda(m)+(1-\theta)\beta_\Lambda(\bar{m}),\end{aligned}$$

thus $\beta_\Lambda(\cdot)$ is convex on $\mathbb{R}_+^K$. □

Proof of Lemma 8.1. We set,

$$f(\epsilon) = \frac{1}{2}\big[\beta^*_{\epsilon,0} - \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n)\big].$$

Since $\beta(\cdot)$ is continuous and $\mathcal{M}(n)$ is compact, by proposition 8.1 we know $f(\epsilon) > 0$. Suppose the result is not true for our specific choice of $f(\epsilon)$ then there exists $\epsilon > 0$ and sequence $\delta_k \to 0$ as $k \to \infty$ such that,

$$\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) + f(\epsilon) > \beta^*_{\epsilon,\delta_k}.$$

By compactness we may take a further subsequence $\{m^{k_l}\}_{l\in\mathbb{N}}$ in $\mathbb{R}_+^K$ with $m^{k_l} \to \bar{m}$ as $l \to \infty$, for some $\bar{m} \in \mathbb{R}_+^K$ such that,

$$|\sum_{j:j\in i} m_{ji}^{k_l} - n_i| \le \delta_{k_l}, \qquad \inf_{m'\in\mathcal{M}(n)} ||m^{k_l} - m'|| \ge \epsilon,$$

and

$$\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) + f(\epsilon) > \beta(m_{ji}^{k_l}).$$

Thus by the continuity of $\beta(\cdot)$,

$$\sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) + f(\epsilon) \ge \beta(\bar{m}),$$

but this contradicts our definition of $f(\epsilon)$ because $\bar{m}$ satisfies $\sum_{j:j\in i} \bar{m}_{ji} = n_i$ $\forall i \in \mathcal{I}$ and $\sup_{m'\in\mathcal{M}(n)} ||\bar{m} - m'|| \ge \epsilon$, and so,

$$\beta(\bar{m}) \ge \beta^*_{\epsilon,0} > \sum_{i:n_i>0} n_i \log \Lambda_i^{PF}(n) + f(\epsilon).$$

□

References

[1] Billingsley, P. (1999). *Convergence of Probability Measures*. Wiley, New York.

[2] Bonald, T. and Massoulie, L. (2001). Impact of Fairness on Internet Performance. *Proc. of ACM Sigmetrics 2001* **29** 82-91.

[3] Bonald, T. and Proutière, A. (2002). Insensitivity in processor-sharing networks. *Performance Evaluation* **49** 193-203.

[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* **30** 89-148.

[6] Chen, H. and Yao, D.D. (2001). *Fundamentals of Queueing Networks: Performance, Asymptotics and Optimization*. Springer, New York.

[7] Dembo, A. and Zeitouni, O. (1998). *Large Deviations Techniques and Applications*. Springer, New York.

[8] Doyle, P.G. and Snell, J.L. (1984). Random walks and electric networks. *The Mathematical Association of America*.

[9] De Veciana, G., Lee, T.J. and Konstantopoulos, T. (1999). Stability and performance analysis of networks supporting services with rate control - could the Internet be unstable? *IEEE/ACM Transactions on Networking* **9**, No. 1, 2-14.

[10] 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. *Submitted to Ann. Appl. Probab.*

[11] 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. *Performance Evaluation Review* **35** 36-38.

[12] Kelly, F.P. (1979). *Reversibility and Stochastic Networks*. Wiley, Chicester.

[13] Kelly, F.P. (1989). On a Class of Approximations for Closed Queueing Networks. *Queueing Systems*, **4** 69-76.

[14] Kelly, F.P. (1997). Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications* **8** 33-37.

[15] Kelly, F. P. and Williams, R. J. (2004). Fluid Model For a Network Operating Under A Fair Bandwidth-Sharing Policy. *Ann. Appl. Probab.* **14** 10551083.

[16] Kingman, J. F. C. (1969). Markov population processes. *J. Appl. Prob.*, **6** 1-18.

[17] Massoulie, L. (2007). Structural Properties of Proportional Fairness: Stability and Insensitivity. *Ann. Appl. Probab.* **17** 809-839.

[18] Massoulie, L. and Roberts, J. (1998). Bandwidth Sharing and Admission Control for Elastic Traffic. *Telecommunication Systems* **15** 185-201.

[19] Massoulie, L. and Roberts, J. (1999). Bandwidth sharing: Objectives and algorithms. IEEE Infocom 1999 & *IEEE/ACM Transactions on networking 2002*, **10**, No. 3, 320328.

[20] Miyazawa, M. (1993). Insensitivity and product-form decomposability of reallocatable GSMP. *Adv. in Appl. Probab.* **25** 415-437.

[21] Norris, J.R. (1997). *Markov Chains*. Cambridge University Press.

[22] Pittel, B. (1979). Closed exponential networks of queues with saturation: The Jackson-type stationary distribution and its asymptotic analysis. *Math. Oper. Res* **4** 357-378.

[23] Schassberger, R. (1978). Insensitivity of steady state distributions of generalized semi-Markov processes with speeds. *Adv. Appl. Probab.* **10** 836-851.

[24] Schweitzer, P.J. (1979). Approximate analysis of multiclass closed networks of queues. Proceedings of the international conference on stochastic control and optimization, Free University, Amsterdam.

[25] Shah, D. and Wischik, D.J. (2008). Heavy Traffic Analysis of Optimal Scheduling Algorithms for Switched Networks. *Submitted to Ann. Appl. Probab.*.

[26] Srikant, R. (2004). *The Mathematics of Internet Congestion Control*. Birkhauser.

[27] Whittle, P. (1985). Partial Balance and Insensitivity. *J. Appl. Probab.* **22** 168-176.