



Stochastic Systems

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Flow-level Convergence and Insensitivity for Multi-class Queueing Networks

Neil Stuart Walton

To cite this article:

Neil Stuart Walton (2012) Flow-level Convergence and Insensitivity for Multi-class Queueing Networks. *Stochastic Systems* 2(1):115-148. <https://doi.org/10.1287/11-SSY025>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Stochastic Systems*. Copyright 2012 The Author(s). <https://doi.org/10.1287/11-SSY025>, used under a Creative Commons Attribution License: <http://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2012, The author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

FLOW-LEVEL CONVERGENCE AND INSENSITIVITY FOR MULTI-CLASS QUEUEING NETWORKS

BY NEIL STUART WALTON

University of Amsterdam

We consider a multi-class queueing network as a model of packet transfer in a communication network. We define a second stochastic model as a model of document transfer in a communication network where the documents transferred have a general distribution. We prove the weak convergence of the multi-class queueing process to the document transfer process. Our convergence result allows the comparison of general document size distributions, and consequently, we prove general insensitivity results for the limit queueing process. We discuss how this separation of time-scales method of proving insensitivity may be applied to other insensitive queueing systems.

1. Introduction. Introduced by Roberts and Massoulié [14], bandwidth sharing networks are a popular stochastic model of flow-level document transfer across a communication network. Here documents of different sizes arrive stochastically and are simultaneously transferred at a rate which depends on the number of documents in transfer. Bandwidth sharing networks are a flow-level model of a packet-switched network. In practice, documents are broken into smaller packets which are then sequentially transferred across a queueing network. Bandwidth networks abstract away these packet-level dynamics and instead documents are assumed to be continuously transferred.

Bandwidth sharing networks can have desirable properties. In particular, Bonald and Proutiere [4] classified the set of bandwidth sharing networks which have a certain *insensitivity* property. Insensitive networks have a stationary distribution which only depends on the expected size of documents in transfer and, thus, are tractable for more general document types than idealized models with exponentially distributed document sizes, which are more typically analysed. In particular, Bonald and Proutiere [4] presented an insensitive bandwidth network which can be seen as the flow-level analogue of a network of multi-class queues. This particular model is our principle object of interest.

Received February 2011.

AMS 2000 subject classifications: Primary 60K25, 60G17; secondary 90B18.

There has been numerous works studying fluid and heavy traffic limits of bandwidth sharing networks with general document sizes [7, 8, 11, 13]. However, there is currently less work devoted to proving that bandwidth sharing networks are the limit of packet-switching networks. This motivates the two main results of this paper:

- Theorem 5.1: We use a separation of time-scales argument to prove that a bandwidth sharing network with general document sizes is the limit of a sequence of packet switched networks.
- Corollary 6.1: We prove the insensitivity of the limit bandwidth sharing network.

Such separation of time-scale limits can provide an effective mechanism to describe and analyse stochastic networks, for instance, see the recent work of Feuillet and Robert [6]. In the context of Bandwidth sharing, in [19], a separation of time-scales proof was used to prove convergence of a packet-switch network to a bandwidth sharing network with exponentially distributed document sizes. We extend this result to a much more general class of models each of which allows for a general distribution of documents sizes. An appealing feature of the convergence proof is that the result can then be used as a method to prove insensitivity of the limit bandwidth sharing network. This gives the two main contributions of this paper:

1. We formally justify that a class of bandwidth sharing networks with general document sizes are the limit of a packet-switched network.
2. We apply separation of time-scales result to a natural prelimit model in order to prove insensitivity of its limit model.

We note that other methods for proving general insensitivity results could be applied to our bandwidth sharing model: the method of stages, Barbour [1]; by coupling of both renewal and Markov processes, Zachary [20]; or by directly verifying stationarity of the Markov processes' forward-equations, Virtamo [18]. However, the separation of time-scales approach is appealing because a general insensitivity result is a natural consequence of convergence from its prelimit model. Shortly, we will elaborate on this as a method of proving insensitivity, but first we briefly describe the convergence result proven in this paper.

1.1. *An informal description of the results.* We consider a multi-class queueing network. The queueing network processes documents along different routes. Arriving documents are divided into packets which are sent across the network one-by-one. Once all the packets in a document are sent the document departs.

We could describe the state of this queueing network in several ways: we could consider the explicit behaviour of the network, Q , by storing residual document sizes and the location of all packets on their routes; or, we could consider the positions of packets only \tilde{Q} , in doing so, we would ignore information about the residual sizes of documents; or, finally, we could consider the residual document sizes only Y , and thus ignore precise information about the positions of packets on their routes. These descriptions form an explicit description, a packet-level description and a flow-level description of our network.

We are interested in the interactions of a queueing network at these levels. In particular, the rate at which packets are transferred in a modern communication network is often an order of magnitude larger than the time it takes to transfer a document. Thus, given the number of documents in transfer, we should be able to abstract away the packet-level behaviour of the network. More formally, conditional on the number of documents in transfer, the quick transition of packets within the network should imply that the distribution of packets within the network converges quickly to its stationary distribution, and thus the processing of documents is best described by the stationary behaviour of the packet-level queueing network. We provide a mathematical justification of this phenomenon.

We do so by taking a sequence of multi-class queueing networks $Q^{(c)}$. Along this sequence, we increase the rate that queues process packets by a factor c and accordingly increase the sizes of documents by a factor c . In this regime, packet transitions occur on a time scale of order $O(1/c)$, whilst document transfers remain at a time-scale of $O(1)$. Thus in this limit, between document arrival-departure events, the packet-level state \tilde{Q} will converge to its stationary distribution and document transfers will receive a constant rate of transfer given by the stationary throughput of the packet-level network. Figure 1 below gives a scaling of such queueing network.

In this paper, we show that in this limit, if we let the initial state and document sizes converge, then the times at which documents arrive and depart the network will converge. The formal result, Theorem 5.1 proves that the flow-level state of queueing networks converges in the Skorohod topology to a flow-level model of the network.

With this result, we can consequently prove a general insensitivity result for these flow-level networks. In this context, insensitivity means that the stationary distribution of number of documents in transfer only depends on the mean size of documents. The Skorohod convergence implies that the stationary distribution for the pre-limit network converges to the stationary distribution for the limit network. Consequently in Corollary 6.1, we demon-




FIG 1. The picture above gives the transition rates for a multi-class queueing network. Queues process packets at rates given by $c\phi_1(m_1)$ and $c\phi_2(m_2)$. Thus the rate packets are transferred between queues is of order $O(c)$. Documents on routes 0, 1 and 2 arrive as a Poisson processes of rates ν_0, ν_1, ν_2 . Therefore documents arrive at a rate of order $O(1)$. In this example, documents on routes 0, 1 and 2 have a geometric distribution with parameters μ_0/c , μ_1/c and μ_2/c , respectively. Now consider the rate documents depart the network. For route 0, for example, the rate documents depart is of the order of $c\phi_2(m_2) \times \mu_0/c = \mu_0\phi_2(m_2)$. Thus document departures occur at a rate of order $O(1)$. This justifies a separation of time-scales between packet transfer and document transfer. This separation of time-scales will be required to form a limit process from a sequence of multi-class queueing networks.

strate that the limit queueing network is insensitive amongst all non-atomic document size distributions.

1.2. *Time-scales and insensitivity: A processor-sharing example.* Although we consider a specific bandwidth sharing model in this paper, the techniques applied here could be applied more broadly as a method to analyse models which are the insensitive limit of simple, fine-grained systems. We motivate this claim with the example of a processor-sharing queue.

A processor-sharing queue divides a unit of service equal amongst all jobs at the queue. Processor-sharing queues are well understood and well known to be insensitive [10, Chapter 3.3]. Indeed, it is such a well accepted model that one often omits the fact that the processor-sharing queue was originally constructed as an idealize limit of a ‘round-robin’ computer processing system, see Klienrock [12]. We can argue that insensitivity can be proven as a consequence of convergence from its natural pre-limit system. The informal arguments we here can be rigorized using the methods applied in this paper.

We consider a round-robin model of a computer processor which processes discrete units of work. Here, jobs arrive at the processor as a Poisson process of rate ν . Each arriving job chooses a position uniformly at random amongst the set of unprocessed jobs. Each arriving job, consists of a number of units

of work that is independent and identically distributed to a random variable W . We suppose it takes an amount of time to process a unit of work that is independent, exponentially distributed with mean one. In order to share its processing capacity, the processor cycles clockwise between the unprocessed jobs as a Poisson process of rate λ . For the sake of tractability, we also suppose after completing all work in a job the processor moves uniformly at random to one of the remaining unprocessed jobs.

We can consider a Markov process that records, x , the current position of the processor and, $w = (w_1, \dots, w_n)$, the number of units of work to be processed from each job at the queue. Using Kelly's Lemma [10, Theorem 1.13], one can verify that the stationary reversed process of this Markov process is a Markov process recording a processor moving anti-clockwise and recording the number of units of work that have so far been processed from each job. As a consequence the stationary distribution of this Markov process is

$$(1) \quad \pi(x, w) = (1 - \nu \mathbb{E}W) \times \frac{1}{n} \times \prod_{k=1}^n \nu \mathbb{P}(W \geq w_k).$$

Note within our Markov process there are two further stochastic processes: one which records the position of the processor, x , and another which records the residual work, $w = (w_1, \dots, w_n)$.

As Klienrock describes [12], there is a separation of time-scales argument where the above model converges to a processor-sharing queue. In particular, if we take a limit where we increase by the same factor the rate of cycling, the rate of processing work and the amount of work in each arriving job, then at any time instant the stochastic process describing the position of the processor will be best described by a uniform distribution on the number of jobs and the stochastic processes describing the residual work will be a processor-sharing queue. Similar to our proof of Theorem 5.1, we could give a formal argument of weak convergence of our pre-limit round-robin queue to a processor-sharing queue, and similar to Corollary 6.1, we could observe that our pre-limit stationary distribution (1) converges to

$$\pi(w)dw_1 \dots dw_n = (1 - \nu \mathbb{E}W) \times \frac{1}{n} \times \prod_{k=1}^n \nu \mathbb{P}(W \geq w_k)dw_1 \dots dw_n.$$

We can then conclude that this must be the stationary distribution of our limit process. Because the stationary number of jobs in this system only depends of mean size of each job, this would then give a proof of insensitivity of the processor-sharing queue.

The convergence result proves the existence of our model from a natural pre-limit model. In contrast to a method introducing mixtures of gamma distributions in order to prove insensitivity, notice, we prove insensitivity as an immediate consequence of the convergence used to justify our process. In other words, insensitivity is inherited from our pre-limit model. Furthermore, in contrast insensitivity proofs which require a well developed theory of continuous-state Markov processes, we use a general method which solely requires convergence from its primitive pre-limit process.

It would seem likely that other examples will fit within this framework. For instance, van de Ven et al. [17] prove phase-type insensitivity of a flow-level the model of a wireless network. Here nodes continuously transfer information and have phase-type back-off periods. This might suggest that that this model can be proven to be insensitive as a limit of the packet-level dynamics of the system.

1.3. *Organisation.* In Section 2, we give some notation used throughout the paper. In Section 3, we review results on some well understood product form queueing networks. Readers familiar with such networks may wish to acquaint themselves with the notation used and then later refer back to the specific results proven in this section. In Section 4, we define the bandwidth sharing networks which will be the limit of our multi-class queueing networks. We, also, define what it means for these networks to be insensitive. In Section 5, we prove the main convergence result Theorem 5.1. Finally, in Section 6, we prove the insensitivity of these queueing networks, Corollary 6.1.

2. Notation and network structure. We let the 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 $i = \{j_1^i, \dots, j_{k_i}^i\} \in \mathcal{I}$, we associate an order $(j_1^i, \dots, j_{k_i}^i)$. We allow for queues to be repeated in our route order. For $i \in \mathcal{I}$ and $j \in i$, we let $\zeta_{ji} \in \mathbb{N}$ be the number of times queue j is included in ordering $(j_1^i, \dots, j_{k_i}^i)$. 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 let $K = |\mathcal{K}|$. We will view a multi-class queueing network model as transferring a number of documents across the different routes of the network. The vector $n = (n_i : i \in \mathcal{I}) \in \mathbb{Z}_+^I$ will denote the number of documents in transfer across the routes of the network. We also let the vector $m = (m_{ji} : (j, i) \in \mathcal{K}) \in \mathbb{Z}_+^K$ refer to the number of packets in transfer across each route at each queue. That is m_{ji} is the number of packets on route i at queue j . We also define the number of packets in transfer at a

queue to be

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

For each $n \in \mathbb{Z}_+^I$, we define $S(n) = \{m \in \mathbb{Z}_+^K : \sum_{j:j \in i} m_{ji} = n_i \forall i \in \mathcal{I}\}$, that is the set of queue sizes with n documents in transfer on each route.

3. Multi-class queueing networks. In this section, we present some known queueing networks that will be studied subsequently. In order to model the transfer of documents across a packet switching network, we define a special case of these queueing networks where packets have a specific routing structure. We then introduce closed queueing networks as described in [10, Section 3.4].

3.1. *Multi-class queue.* First, we define what we will call a *multi-class queue*. We consider a single queue j from a set of queues \mathcal{J} . We call the customers of this queue packets. The queue will receive packet arrivals from different classes. The set of classes will consist of a set of packet route choices \mathcal{C} .¹ Packets occupy different positions within a queue. Given there are $m_j \in \mathbb{Z}_+$ packets at queue j , packets may occupy positions $1, 2, \dots, m_j$. Packets of each route at the queue require an independent exponentially distributed service requirement with mean 1. Given there are $m_j \in \mathbb{Z}_+$ packets at queue j , the total service devoted to packets is given by $\phi_j(m_j)$. We assume $\phi_j(m_j) > 0$ if $m_j > 0$. This service is then divided amongst packets within the queue. Given there are $m_j \in \mathbb{Z}_+$ packets at queue j , a proportion $\gamma_j(l, m_j)$ of service is devoted to the packet in position $l \in \{1, \dots, m_j\}$ of queue j . Since $\gamma_j(\cdot, m_j)$ represents a proportion,

$$\sum_{l=1}^{m_j} \gamma_j(l, m_j) = 1, \quad m_j \in \mathbb{N}.$$

Upon completing its service the packet at position l will leave the queue and the packets at positions $l + 1, \dots, m_j$ will move to positions $l, \dots, m_j - 1$, respectively.

We assume packets of class $c \in \mathcal{C}$ will arrive at the queue from independent Poisson processes of rate ρ_{jc} . Given there are m_j packets at the queue an arriving packet will move to position $l \in \{1, \dots, m_j + 1\}$ with probability $\delta_j(l, m_j + 1)$. Once again as $\delta_j(\cdot, m_j + 1)$ represents a proportion

$$\sum_{l=1}^{m_j+1} \delta_j(l, m_j + 1) = 1, \quad m_j \in \mathbb{Z}_+.$$

¹Although for now we choose \mathcal{C} to be arbitrary, we will later consider $\mathcal{C} = \{i \in \mathcal{I} : j \in i\}$, the set of routes using queue j .

When a packet arrives at position l the packets in positions l, \dots, m_j will move to positions $l + 1, \dots, m_j + 1$, respectively.

Let $q^j = (c_1^j, \dots, c_{m_j}^j) \in \mathcal{I}^{m_j}$, for $m_j > 0$, give the state of queue j . Let function, $T_{\cdot, (j,l)}^c$ denote the arrival of a class c packet to position l in queue j and let function $T_{(j,l)}^c$ denote the departure of a class c packet in position l . Thus the state of this queue forms a continuous-time Markov chain with transition rates given by,

$$r(q^j, q'^j) = \begin{cases} \rho_c \delta_j(l, m_j + 1) & \text{for } q'^j = T_{\cdot, (j,l)}^c q^j, l = 1, \dots, m_j + 1, c \in \mathcal{C} \\ \phi_j(m_j) \gamma_j(l, m_j) & \text{for } q'^j = T_{(j,l)}^c q^j, c_l^j = c, l = 1, \dots, m_j, \\ 0 & \text{otherwise.} \end{cases}$$

The queue itself will not discriminate between different packet's classes and thus the stationary distribution of the queue size will be oblivious to different packets' route type. Ignoring packet classes, when stationary M_j the Markov chain recording the total number of packets at the queue is reversible. Given m_j , routes of the packets in positions $1, 2, \dots, m_j$ are independent. The probability a packet in a given position is from route i is $\frac{\zeta_{ji} \rho_i}{\sum_{r \ni j} \zeta_{jr} \rho_r}$. Thus letting Markov chain Q_j record the position and routes of packets at queue j and letting $\mathcal{Q}_j = \cup_{m_j=1}^{\infty} \mathcal{I}^{m_j}$ gives all possible states of the queue, we can calculate the stationary distribution of the queue.

PROPOSITION 3.1 (BCMP [2], Kelly [9]). *A stationary multi-class queue is quasi-reversible and the stationary distribution of Q_j must be*

$$(2) \quad \mathbb{P}_j(Q_j = (c(1), \dots, c(m_j))) = \frac{1}{B_j} \prod_{l=1}^{m_j} \frac{\rho_{c(l)}}{\phi_j(l)}, \quad (c(1), \dots, c(m_j)) \in \mathcal{Q}_j.$$

Moreover, for $j \in \mathcal{J}$, the process $(M_{jc} : i \in \mathcal{C})$ giving the number of packets of each route type at queue j , has stationary distribution

$$(3) \quad \mathbb{P}(M_{jc} = m_{jc}, \forall c \in \mathcal{C}) = \frac{1}{B_j} \binom{m_j}{m_{jc} : i \in \mathcal{C}} \frac{\prod_{c \in \mathcal{C}} (\rho_c)^{m_{jc}}}{\prod_{l=1}^{m_j} \phi_j(l)}$$

$\forall (m_{jc} : c \in \mathcal{C}) \in \mathbb{Z}_+^{| \mathcal{C} |}$, where we define

$$(4) \quad \binom{m_j}{m_{jc} : c \in \mathcal{C}} = \frac{m_j!}{\prod_{c \in \mathcal{C}} (m_{jc}!)}.$$

The combinatorial term in (4) is required as the probability distribution (3) ignores the order of packets within the queue.

3.2. *Multi-class queueing networks.* A multi-class queueing network with spinning (MQNwS) is a multi-class network of quasi-reversible queues with the following class routing structure. The class of a packet is of the form $c = (i, k, y)$ where $i \in \mathcal{I}$ records the route the packet is on, $k \in \mathbb{N}$ records the stage of the packet on route i and $y \in \mathbb{N}$ records the packet's residual document size, that is the remaining number of times the packet must traverse its route. Routing through classes occurs in the following way. If route i has associated route order $(j_1^i, \dots, j_{k_i}^i)$ then as a Poisson process of rate $\nu_i \mathbb{P}(X_i = x)$ class $(i, 1, x)$ packets arrive at queue j_1^i . Here X_i is a random variable with values in \mathbb{N} and with mean $\mu_i^{-1} < \infty$. For $k = 1, \dots, k_i - 1$, a class (i, k, y) packet on departing queue j_k^i will join queue j_{k+1}^i and become a class $(i, k + 1, y)$ packet. For a packet that has completed service at the final queue on route $i \in \mathcal{I}$ and has not been fully processed through the network, that is a packet of class (i, k_i, y) with $y > 1$, the packet will join queue j_1^i as a class $(i, 1, y - 1)$ packet. For a route $i \in \mathcal{I}$ packet that has completed its service at the final queue k_i and has been fully processed through the network, that is of class $(i, k_i, 1)$, the packet will depart the network. In addition, we let the constant $\zeta_{ji} \in \mathbb{Z}_+$ give the number of times a packet visits queue j each time it traverses route i . Finally, we define traffic intensities $\rho_i = \frac{\nu_i}{\mu_i}$ for each $i \in \mathcal{I}$.

We can interpret this routing structure in two ways. First, we could consider each packet on route i to arrive as a Poisson process and to repeat its route a number of times that is independent and with distribution equal to X_i . This interpretation leads us to think of a packet as spinning around its route a random number of times. Second, we could consider the network to be transferring documents. Documents which require to be transferred across route i arrive as a Poisson process as of rate ν_i . Each document consists of a number of packets, that is independent and with distribution equal to X_i . These packets are then sent across the network one-by-one until the document is transferred.

For a Markov process description of a MQNwS we record its *explicit state*: we let $q = (q_j : j \in \mathcal{J})$, where $q_j = (c_j(1), \dots, c_j(m_j))$ gives the class of each customer in each occupied position in queue j . Here the class $c_j(l) = (i_j(l), k_j(l), y_j(l))$ records the route, stage and residual document size associated with the l -th packet in queue j . We let \mathcal{Q} define the set of all possible states for this explicit description of our queueing network.

We recall that m_{ji} is the number of route i packets in transfer at queue j and that n_i is the number of route i documents in transfer. As each document

has one packet in transfer in the network at any point in time

$$n_i = \sum_{j \in \mathcal{J}} m_{ji}, \quad i \in \mathcal{I}.$$

We define two further descriptions of the state of a MQNwS: the packet level state and the flow level state.

We define the *packet level state* of a multi-class queueing network with spinning to be, $\tilde{q} = (\tilde{q}_j : j \in \mathcal{J})$, where $\tilde{q}_j = (\tilde{c}_j(1), \dots, \tilde{c}_j(m_j))$ and where $\tilde{c}_j(l) = (i_j(l), k_j(l))$ records the route and stage associated with the l -th packet in queue j . We let $\tilde{\mathcal{Q}}$ define the set of all possible packet level states for this description of our queueing network. The packet level state of a MQNwS is concerned with the position and route of packets but not of the state of document transfer. Similarly the flow level state is interested in the state of document transfer and not in the specific position of packets.

We define the *flow level state* of a MQNwS to be, given by vector

$$y = (y_{ik} : k = 1, \dots, n_i, i \in \mathcal{I})$$

Here, we order elements so that $y_{ik} \leq y_{ik+1}$ for all $k = 1, \dots, n_i - 1$. Note, we record no information about each packet's position on its route. As described above n_i refers to the number of route i documents in transfer on route i and k indexes each specific packet in transfer on route i . The number y_{ik} is the residual document size of the k -th document in route i , that is the number of packets yet to be transferred from the document. We let \mathcal{Y} be the set of flow level states achievable by a MQNwS.

The processes associated with the packet level or flow level state of a multi-class queueing network with spinning need not be Markov. However, these state descriptions will be useful for proving weak convergence results. For this purpose we define on \mathcal{Y} the norm

$$(5) \quad \|y - y'\| = \begin{cases} \max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i} |y_{ik} - y'_{ik}| & \text{if } n_i = n'_i, \quad \forall i \in \mathcal{I}, \\ \infty & \text{otherwise.} \end{cases}$$

3.2.1. Stationary behaviour. We now calculate certain quantities associated with the stationary distribution of a MQNwS. As a direct consequence of known reversibility results [10, Theorem 3.1], we can calculate the stationary distribution of a MQNwS.

THEOREM 3.1. *The explicit state of an ergodic multi-class queueing network with spinning has stationary distribution,*

$$(6) \quad \mathbb{P}(Q = q) = \prod_{j \in \mathcal{J}} \frac{1}{B_j} \prod_{l=1}^{m_j} \left(\frac{\nu_{i_j(l)} \mathbb{P}(X_{i_j(l)} \geq y_j(l))}{\phi_j(l)} \right), \quad q \in \mathcal{Q},$$

provided

$$(7) \quad B_j = \sum_{m_j=1}^{\infty} \left(\prod_{l=1}^{m_j} \frac{\sum_{i:j \in i} \zeta_{ji} \rho_i}{\phi_j(l)} \right) < \infty, \quad \forall j \in \mathcal{J}.$$

PROOF. A multi-class queueing network with spinning is a network of quasi-reversible queues with a deterministic routing structure. It is known, [10, Theorem 3.1], that a network of quasi-reversible queues has a stationary distribution

$$\mathbb{P}(Q = q) = \prod_{j \in \mathcal{J}} \mathbb{P}(Q_j = q_j)$$

where,

$$\mathbb{P}(Q_j = q_j) = \frac{1}{B_j} \prod_{l=1}^{m_j} \frac{\beta_{jc_j(l)}}{\phi_j(l)}, \quad j \in \mathcal{J},$$

and where β_{jc} solves the traffic equations

$$\beta_{jc} = \nu_{jc} + \sum_{l,d} \beta_{l,d} p_{ld,jc}, \quad j \in \mathcal{J}, c \in \mathcal{C}.$$

Here ν_{jc} is the arrival rate of class c customers at queue j and $p_{jc,ld}$ gives the packet routing probabilities, which in our case are, for $c = (i, k, y)$,

$$p_{jc,ld} = \begin{cases} 1 & \text{if } k < k_i, d = (i, k + 1, y) \text{ and } l = j_{k+1}^i, \\ 1 & \text{if } k = k_i, d = (i, 1, y - 1), y > 0 \text{ and } l = j_1^i, \\ 1 & \text{if } y = 1, k = k_i, (l, d) = \cdot, \\ 0 & \text{otherwise.} \end{cases}$$

In this way, packets are transferred between queues, the next packet is injected at the ingress and document departures occur.

So, all that is needed is to verify that $\tilde{\beta}_{j,c}$ solves the traffic equations along our deterministic path. Observe that, for $k > 1$, $\tilde{\beta}_{j,(i,k,y)} = \tilde{\beta}_{j,(i,k-1,y)}$ and, for $k = 1$

$$\begin{aligned} \tilde{\beta}_{j,(i,1,y)} &= \nu_i \mathbb{P}(X_i \geq y_i) = \nu_i \mathbb{P}(X_i = y) + \nu_i \mathbb{P}(X_i \geq y - 1) \\ &= \nu_{i,(i,1,y)} + \tilde{\beta}_{j,(i,k_i,y+1)}. \end{aligned}$$

This verifies the traffic equations are satisfied and hence gives the result. \square

The condition (7) is the necessary and sufficient for a multi-class queueing network with spinning to be ergodic and thus is equivalent to the assumptions ergodicity in subsequent results. We encapsulate this in the following assumption:

ASSUMPTION 1. *Unless stated otherwise we assume a multi-class queueing network with spinning satisfies the following necessary and sufficient condition for ergodicity:*

$$(8) \quad B_j = \sum_{m_j=1}^{\infty} \left(\prod_{l=1}^{m_j} \frac{\sum_{i:j \in i} \zeta_{ji} \rho_i}{\phi_j(l)} \right) < \infty, \quad \forall j \in \mathcal{J}.$$

The following three corollaries are a consequence of Theorem 3.1. Each of these results require summing over an appropriate set of states. For example, from Theorem 3.1, we can calculate the stationary distribution of the number of packets in transfer along each route at each queue.

COROLLARY 3.1. *Given the stability condition, Assumption 1, $M = (M_{ji} : (j, i) \in \mathcal{K})$, the number of packets in transfer across each route at each queue, has stationary distribution,*

$$(9) \quad \mathbb{P}(M = m) = \prod_{j \in \mathcal{J}} \left(\frac{1}{B_j} \binom{m_j}{m_{ji} : i \ni j} \frac{\prod_{i:j \in i} (\zeta_{ji} \rho_i)^{m_{ji}}}{\prod_{l=1}^{m_j} \phi_j(l)} \right)$$

PROOF. We know from Theorem 3.1 that our queueing network has a stationary distribution equal to that of a simpler queueing network. In this simpler queueing network, each queue j behaves independently in isolation and where class $c = (i, k, y)$ packets, with $j_k^i = j$, arrive at queue j as a Poisson process of rate $\nu_i \mathbb{P}(X_i \geq y)$. Let us work with this simpler but equivalent queueing model. In this model, route i packets arrive into queue j as a Poisson process of rate

$$\sum_{k:j_k^i=j} \sum_{y=1}^{\infty} \nu_i \mathbb{P}(X_i \geq y) = \zeta_{ji} \rho_i.$$

So queue j will have an independent stationary distribution exactly of the form of (2). So, as in Section 3.1, by ignoring packet positions we can gain distribution (3) for each queue. Thus by independence equation (9) holds for the network. \square

REMARK 1. Observe that distribution (9) only depends on the distribution of X_i through its mean $\frac{1}{\mu_i}$. Thus the stationary distribution of a MQNwS only depends on the distribution of document sizes through their mean size. This suggests a form of insensitivity holds. This point is noted by Massoulié and in the thesis of Proutiere [16]. Similar observations are made earlier in [10] for individual queues. We will discuss this observation in more detail in Chapter 3.

We can express the stationary distribution of the number of documents in transfer on each route, $N = (N_i : i \in \mathcal{I})$. We define $\mathcal{Q}(n)$ and $\tilde{\mathcal{Q}}(n)$ to be the set of explicit states and packet level states a MQNwS that occur with positive probability given there are $n \in \mathbb{Z}_+^I$ documents in transfer on each route. We also let $S(n) = \{m \in \mathbb{Z}_+^K : \sum_{j:j \in i} m_{ji} = n_i, \forall i \in \mathcal{I}\}$ be the set of route-queue states achievable given there are $n \in \mathbb{Z}_+^I$ documents in transfer.

COROLLARY 3.2. *Given the stability condition Assumption 1 is satisfied. For a MQNwS, $N = (N_i : i \in \mathcal{I})$ the number of documents in transfer on each route has stationary distribution*

$$(10) \quad \mathbb{P}(N = n) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \rho_i^{n_i}, \quad n \in \mathbb{Z}_+^I,$$

where we define

$$(11) \quad B := \prod_{j \in \mathcal{J}} B_j,$$

$$(12) \quad B_n := \sum_{m \in S(n)} \prod_{j \in \mathcal{J}} \binom{m_j}{m_{ji} : i \ni j} \left(\frac{\prod_{i:j \in i} \zeta_j^{m_{ji}}}{\prod_{l=1}^{m_j} \phi_j(l)} \right), \quad n \in \mathbb{Z}_+^I.$$

Finally, we give the stationary distribution for the packet level state of a MQNwS, $\tilde{Q} = (\tilde{Q}_j : j \in \mathcal{J})$. And, we also give the stationary distribution the flow level state of a MQNwS, $Y = (Y_{ik} : k = 1, \dots, N_i, i \in \mathcal{I})$.

COROLLARY 3.3. *Given the stability condition Assumption 1 is satisfied, the stationary packet level state of a MQNwS, $\tilde{Q} = (\tilde{Q}_j : j \in \mathcal{J})$, has distribution*

$$\mathbb{P}(\tilde{Q} = \tilde{q}) = \prod_{j \in \mathcal{J}} \frac{1}{B_j} \prod_{l=1}^{m_j} \frac{\rho_{i_j(l)}}{\phi_j(l)}, \quad \tilde{q} \in \tilde{\mathcal{Q}}.$$

The stationary flow level state of a MQNwS, $Y = (Y_{ik} : k = 1, \dots, N_i, i \in \mathcal{I})$, has distribution

$$(13) \quad \mathbb{P}(Y = y) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \binom{n_i}{n_{iy} : y \in \mathbb{N}} \prod_{y \in \mathbb{N}} (\nu_i \mathbb{P}(X_i \geq y_{ik}))^{n_{iy}}, \quad y \in \mathcal{Y}.$$

Here n_{iy} is the number of route i packets with residual file size y . Also we define,

$$\binom{n_i}{n_{iy} : y \in \mathbb{N}} := \frac{n_i!}{\prod_{y \in \mathbb{N}} (n_{iy}!)}.$$

The above two corollaries simply involve summing distribution (6) over the specified set of states. We omit the explicit calculation in their proof.

3.3. Closed multi-class queueing network. A closed multi-class queueing network behaves as an MQNwS except that document arrivals and departures are forbidden, see [10, Section 3.4]. In effect the network behaves as if there are a fixed number of infinitely large documents in transfer. We now more formally define a closed multi-class queueing network.

Given there are $n \in \mathbb{Z}_+^I$ packets on each route, a closed multi-class queueing network is a packet level Markov process on the states $\tilde{Q}(n)$. We now define the class and routing structure of this queueing network. The class of a packet is of the form $\tilde{c} = (i, k)$ where $i \in \mathcal{I}$ records the route a packet is on and $k \in \mathbb{N}$ records the stage of the packet on its route i . Routing through classes occurs in the following way. For $k = 1, \dots, k_i - 1$, a class (i, k) packet on departing queue j_k^i will join queue j_{k+1}^i and become a class $(i, k + 1)$ packet. A class (i, k_i) packet that has completed service will join queue j_1^i as a class $(i, 1)$ packet.

This description is sufficient to give a Markov chain description of a closed multi-class queueing network, but is not sufficient for this Markov chain to be irreducible. For example, a network consisting of a single last-come-first-served queue would be reducible. For this reason, we require the following assumption to hold throughout this paper.

ASSUMPTION 2. *We assume for all closed queueing networks in this article that the set of states $\tilde{Q}(n)$ is irreducible.*

It is worth noting that if Assumption 2 is broken then there need not be a unique stationary distribution or a unique stationary throughput for the closed queueing network. Note due to the finite state space of these Markov chains we do not require any stability condition to hold.

As is proven in Section 3.4 of [10], we now give the stationary distribution for this queueing network.

COROLLARY 3.4. *Given Assumption 2, for a closed multi-class queueing network with $n \in \mathbb{Z}_+^I$ documents in transfer across routes, the number of packets in transfer on each route at each queue has stationary distribution*

$$(14) \quad \mathbb{P}_n(M = m) = \frac{1}{B_n} \prod_{j \in \mathcal{J}} \left(\binom{m_j}{m_{ji} : i \ni j} \left(\frac{\prod_{i: j \in i} \zeta_{ji}^{m_{ji}}}{\prod_{k=1}^{m_j} \phi_j(k)} \right) \right),$$

for each $m \in S(n)$, where B_n is defined by (12).

Finally, we can characterise the stationary throughput of packets in a closed multi-class queueing networks.

COROLLARY 3.5. *Given Assumption 2, for a closed multi-class queueing network with $n \in \mathbb{Z}_+^I$ documents in transfer across routes and with $n_i > 0$, the stationary throughput of each route i packet, at stage k and at queue $j = j_k^i$ is*

$$\frac{1}{n_i} \frac{B_{n-e_i}}{B_n},$$

where B_n is defined by (12) and e_i is the i -th unit vector in \mathbb{R}_+^I .

PROOF. The probability the network is in state $m \in \mathbb{Z}_+^K$ is given by (14). Given the network is in state m , by Corollary 3.4 of [10] or from stationary distribution (14), the probability at queue j the packet position $k' \in \{1, \dots, m_j\}$ is traversing route i at stage k is $\frac{1}{\zeta_{ji}} \frac{m_{ji}}{m_j}$. The throughput of the packet in position k' of queue j is $\gamma_j(k', m_j) \phi_j(m_j)$. By our irreducibility assumption, all arrangements of the n_i route i packets are equally likely. Thus the probability this packet is any specific route i packet is $\frac{(n_i-1)!}{n_i!} = \frac{1}{n_i}$. Thus, the stationary throughput of this route i packet is

$$\begin{aligned} & \sum_{\substack{m \in S(n): \\ m_j > 0}} \sum_{k'=1}^{m_j} \gamma_j(k', m_j) \phi_j(m_j) \frac{1}{\zeta_{ji}} \frac{m_{ji}}{m_j} \frac{1}{n_i} \frac{1}{B_n} \prod_{l \in \mathcal{J}} \left(\binom{m_l}{m_{lr} : r \ni l} \right) \left(\frac{\prod_{r:l \in r} \zeta_{lr}^{m_{lr}}}{\prod_{c=1}^{m_l} \phi_l(c)} \right) \\ &= \sum_{\substack{m \in S(n): \\ m_j > 0}} \phi_j(m_j) \frac{1}{\zeta_{ji}} \frac{m_{ji}}{m_j} \frac{1}{n_i} \frac{1}{B_n} \prod_{l \in \mathcal{J}} \left(\binom{m_l}{m_{lr} : r \ni l} \right) \left(\frac{\prod_{r:l \in r} \zeta_{lr}^{m_{lr}}}{\prod_{c=1}^{m_l} \phi_l(c)} \right) \\ &= \sum_{m' \in S(n-e_i)} \frac{1}{n_i} \frac{1}{B_n} \prod_{l \in \mathcal{J}} \left(\binom{m'_l}{m'_{lr} : r \ni l} \right) \left(\frac{\prod_{r:l \in r} \zeta_{lr}^{m'_{lr}}}{\prod_{c=1}^{m'_l} \phi_l(c)} \right) = \frac{1}{n_i} \frac{B_{n-e_i}}{B_n}. \end{aligned}$$

In the first inequality, we used the fact that $\sum_{l=1}^{m_j} \gamma_j(l, m_j) = 1$, $\forall m_j \in \mathbb{N}$. In the second equality, we cancelled terms and substituted $m'_{lr} = m_{lr} - 1$ if $(l, r) = (j, i)$ and $m'_{lr} = m_{lr}$ otherwise. \square

In subsequent chapters, an important quantity will be

$$\Lambda_i^{SN}(n) = \frac{B_{n-e_i}}{B_n},$$

the stationary rate packets are transferred on route i of a closed multi-class queueing network.

4. Bandwidth sharing networks. In this section, we consider a flow level bandwidth sharing model introduced by Massoulié and Roberts [14, 15]. We call these models *Bandwidth Sharing Networks* (BSN). BSNs model the dynamic, elastic transfer rate received by document transfers in a communication network. Multi-class queueing networks with spinning (MQNwS) model packet level dynamics BSNs model document level dynamics. We think of MQNwSs as a microscopic model of a communication network. We think of BSNs as a macroscopic model of a communication network. We will formally relate MQNwS and BSNs.

Massoulié and Roberts [15] discuss the separation of time scales between a certain BSN and MQNwS. In the next section, we will give a proof that a BSN is the limit of a sequence of MQNwS, and thus, we formally justify a separation of time scales. We call our limit flow level model a “spinning network”. The models of this type are considered by Bonald and Proutiere [4] under the name the “Store-Forward Network”.

In performing this analysis, we are able to prove general insensitivity results for the spinning network. As cited by Proutiere [16, Section 3.4] the spinning network was first considered by Massoulié because of its insensitivity. Bonald and Proutiere [4] proved insensitivity for spinning networks with documents with size given by phase type distributions.

In this section, we introduce the bandwidth sharing networks and we define the spinning network. In the next section, and specifically in Theorem 5.1, we prove the main result of this chapter: the weak convergence of a sequence multi-class queueing networks to its spinning network. In Section 6 and specifically in Corollary 6.1, we prove insensitivity results which hold as a consequence of Theorem 5.1.

4.1. *Bandwidth allocations and bandwidth sharing networks.* A *bandwidth allocation policy* is a map $\Lambda : \mathbb{Z}_+^I \rightarrow \mathbb{R}_+^I$. For $n \in \mathbb{Z}_+^I$, the vector $\Lambda(n) = (\Lambda_i(n) : i \in \mathcal{I})$ is a *bandwidth allocation*. Here $\Lambda_i(n)$ represents the rate that route i documents are transferred through each route of a communication network, given there are $n = (n_i : i \in \mathcal{I})$ documents in transfer on each route.

The stochastic model we describe represents the randomly varying number of document transfers within a network. The model is studied as a flow level model of Internet congestion control. We first assume that documents have a size that is exponentially distributed. We will then generalise to document sizes that are independent and of a general distribution.

For document sizes that are independent exponentially distributed, a *bandwidth sharing network operating under bandwidth allocation policy* Λ

(BSN) is a continuous-time Markov chain on \mathbb{Z}_+^I with rates

$$(15) \quad q(n, n') = \begin{cases} \nu_i & \text{if } n' = n + e_i, \\ \mu_i \Lambda_i(n) & \text{if } n' = n - e_i \text{ and } n_i > 0, \\ 0 & \text{otherwise,} \end{cases}$$

for $n, n' \in \mathbb{Z}_+^I$, where e_i is the i -th unit vector in \mathbb{Z}_+^I .

This model can be interpreted as follows: documents wishing to be transferred across route i arrive as a Poisson process of rate ν_i . These documents 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.

The key distinction between this model of document transfer and our previous queueing models is that we do not consider packet level dynamics. These dynamics are abstracted away, and instead, we only consider the flow-level descriptions of the network's state.

We can generalise BSNs to allow the transfer of documents with an independent arbitrarily distributed size. In this case, similar to the flow level state of a multi-class queueing network with spinning, we record the *flow level state* of a generalised BSN. For each document in transfer, we will record the documents residual size, that is the amount of the document that is still to be processed. Given there are $n = (n_i : i \in \mathcal{I})$ documents in transfer, the flow level state of a generalised bandwidth sharing network is given by the vector

$$(16) \quad y = (y_{ik} : k = 1, \dots, n_i, i \in \mathcal{I}).$$

Here $y_{ik} \in (0, \infty)$ is the residual document size of the k -th document in transfer on route i . We order elements so that $y_{ik} \leq y_{i,k+1}$ for $k = 1, \dots, n_i - 1$.

The dynamics of this generalised BSN are morally the same as our previous definition: documents arrive as a Poisson process; documents are transferred at an elastic rate depending on the number of documents in transfer along different routes and documents depart once transferred.

More explicitly, the dynamics of this model are defined as follows. Documents for transfer on route i arrive as a Poisson process of rate ν_i . An arriving document on route i will then have a residual document size X'_i

added to the flow level description (16). We assume X'_i is an independent positive random variable with finite mean μ_i^{-1} , and we assume X'_i is equal in distribution to some positive random variable X_i . In between a document arrival or departure event, the residual document size of a route i document decreases linearly at rate $\frac{\Lambda_i(n)}{n_i}$. A document on route i departs the network at the instant its residual document size equals 0. At this point, the corresponding document is removed from the network's flow level state description.

Given the network's state y , all future events are a function of y and independent random variables, thus the state description describes this process as a Markov process. As described in Section 3.2, we let \mathcal{Y} be the set of flow level states.

4.2. *Insensitive bandwidth sharing networks.* A bandwidth sharing network, as described above, has stationary distribution π_Y when

$$\pi_Y(A) = \mathbb{P}(Y(0) \in A) \quad \text{implies} \quad \pi_Y(A) = \mathbb{P}(Y(t) \in A),$$

$\forall A \in \mathcal{B}(\mathcal{Y})$ and $\forall t \in \mathbb{R}_+$. Here $\mathcal{B}(\mathcal{Y})$ is the Borel σ -field defined on the set of flow level states \mathcal{Y} from norm (5). We say that a random variable X with values in \mathbb{R}_+ is *non-atomic* if $\mathbb{P}(X = x) = 0$ for all $x \in \mathbb{R}_+$.

We say that a bandwidth sharing network is *insensitive* to non-atomic distributions with stationary distribution $\pi_N = (\pi_N(n) : n \in \mathbb{Z}_+^I)$, if every generalised BSN with non-atomic document size distributions, mean document sizes $(\frac{1}{\mu_i} : i \in \mathcal{I})$ has a stationary distribution π_Y satisfying

$$\pi_N(n) = \mathbb{P}_{\pi_Y}(N(0) = n), \quad \forall n \in \mathbb{Z}_+^I.$$

In other words, the distribution π_N only depends on the document size distribution through its mean $(\mu_i^{-1} : i \in \mathcal{I})$.

4.3. *Spinning networks.* Bandwidth allocations represent the stationary rate of document transfer, given the number of documents in transfer on each route. From Corollary 3.5, we can define a bandwidth allocation that represents the stationary behaviour of a MQNwS. We define a *spinning allocation* to be the stationary throughput of a closed multi-class queueing network. That is for each $\forall n \in \mathbb{Z}_+^I$, we define

$$(17) \quad \Lambda_i^{SN}(n) := \frac{B_{n-e_i}}{B_n},$$

where e_i is the i -th unit vector in \mathbb{R}_+^I and B_n is defined by (12). The bandwidth sharing network defined by a spinning allocation policy is called a

spinning network. Proutiere [16] notes that this bandwidth allocation is first defined by Laurent Massoulié. Insensitivity results on this bandwidth allocation are explored in Bonald and Proutiere[4].

5. Convergence of open queueing networks to spinning networks.

We are now in a position to prove the main results of this paper. The bandwidth sharing networks of [14] are intended to represent the flow level dynamics of document transfer in a packet switched network. The aim of this section is to formally justify this interpretation for spinning networks. As a consequence of this analysis, we are able to formally prove insensitivity of spinning networks.

For exponential document sizes and processor sharing queues of fixed capacity, it has been demonstrated that a series of multi-class queueing networks converged weakly to the spinning network in the Skorohod topology [19]. In this section, we generalize these arguments to include general document size distributions and for the general queueing networks discussed in Section 3.2. Although our proof is applied to networks of quasi-reversible queues, the proof applied is phrased so that a more diverse range of queueing processes could be considered. In this sense we generalize Theorem 3.1 [19], whose proof is specific to the specific queueing and document sizes considered.

5.1. *Limit and prelimit parameters.* For our limit model, we consider the bandwidth sharing network for the spinning network. We assume documents have a general positive distribution. As discussed in Section 4.1, we assume documents for transfer on route $i \in \mathcal{I}$ have a distribution given by positive random variable $X_i^{(\infty)}$, with finite mean μ_i^{-1} . We let process $Y^{(\infty)} = (Y^{(\infty)}(t) \in \mathcal{Q} : t \in \mathbb{R}_+)$ give the flow level state of this generalised bandwidth sharing network and we let $N^{(\infty)} = (N^{(\infty)}(t) \in \mathbb{Z}_+^{\mathcal{I}} : t \in \mathbb{R}_+)$ give the number of documents in transfer on each route of the spinning network.

For our prelimit model, we consider a sequence of multi-class queueing networks with spinning indexed by $c \in \mathbb{N}$. For this sequence, we assume that the parameters for queues \mathcal{J} , routes \mathcal{I} , route orders $(j_1^i, \dots, j_{k_i}^i)$ and Poisson arrival rates $\nu = (\nu_i : i \in \mathcal{I})$ are all fixed and coincide with the same parameters used previously to define a spinning network. In our sequence of multi-class queueing networks with spinning, we choose to vary the number of packets in each document and the rate at which packets are transferred through the network. For the c -th multi-class queueing network, we let route

i document's size have a distribution $X_i^{(c)}$ such that

$$\frac{X_i^{(c)}}{c} \Rightarrow X_i^{(\infty)}, \quad \text{as } c \rightarrow \infty, \quad i \in \mathcal{I},$$

and we vary the queueing capacities so that $\phi_j^{(c)}(\cdot) = c\phi_j(\cdot)$, for $j \in \mathcal{J}$.

Our choice of scalings are purposefully chosen so that transitions between queues occur at a time scale of order $O(\frac{1}{c})$ and thus the number of transitions before a document departure is of order $O(1)$. See Figure 1 for further explanation.

For $c \in \mathbb{N}$, we let process $Q^{(c)} = (Q^{(c)}(t) \in \mathcal{Q} : t \in \mathbb{R}_+)$ give the explicit queueing description of the c -th multi-class queueing network with spinning and we let process $N^{(c)} = (N^{(c)}(t) \in \mathbb{Z}_+^I : t \in \mathbb{R}_+)$ give the number of documents in transfer on each route of the c -th network. Let $Y^{(c)} = (Y^{(c)}(t) : t \in \mathbb{R}_+)$ and $\tilde{Q}^{(c)} = (\tilde{Q}^{(c)}(t) : t \in \mathbb{R}_+)$ be the respective processes corresponding to the flow level state and packet level state of the c -th multi-class queueing network with spinning.

Associated with the multi-class queueing network with spinning $Q^{(1)}$, we will consider $\bar{Q}_n = (\bar{Q}_n(t) : t \in \mathbb{R}_+)$ the closed multi-class queueing network with $n \in \mathbb{Z}_+^I$ packets on each route. We make the following assumption about each \bar{Q}_n

ASSUMPTION 3. *We assume Assumption 2 holds for \bar{Q}_n for all $n \in \mathbb{Z}_+^I$. That is \bar{Q}_n is an irreducible Markov chain for all $n \in \mathbb{Z}_+^I$.*

As noted in Section 3.3, this assumption excludes reducibility issues which can occur in closed queueing networks where a queue serves a single deterministically chosen packet.

We will also require an assumption on the spinning network $Y^{(\infty)}$.

ASSUMPTION 4. *We assume for $Y^{(\infty)}$ that, almost surely, there are no simultaneous document arrival-departure events.*

This assumption avoids complications associated with the definition of convergence in the Skorohod Topology. Later we will verify that if distribution $X_i^{(\infty)}$ is non-atomic $\forall i \in \mathcal{I}$ then Assumption 4 holds.²

Our main theorem, Theorem 5.1, considers weak convergence on bounded time intervals. Thus, we do not require assumptions on the networks long

²Given the Poisson arrival process of this model Assumption 4 should hold for all BSNs, provided the initial distribution is chosen so that documents do not arrive or depart at the same time.

run behaviour, such as Assumption 1, however we will subsequently require some assumptions for results on insensitivity.

5.2. *Theorem and proof.* We now state and prove the main result.

THEOREM 5.1. *For $c \in \mathbb{N}$, take an multi-class queueing network with spinning $Q^{(c)}$, as described above. Assume Assumptions 3 and 4 hold for each $c \in \mathbb{N}$. Let $Y^{(\infty)}$ denote the flow level state of the spinning network, as described above. If the initial flow level state converges,*

$$(18) \quad \frac{Y^{(c)}(0)}{c} \Rightarrow Y^{(\infty)}(0) \quad \text{as } c \rightarrow \infty$$

then, for each $T > 0$, the stochastic processes converge in the Skorohod topology on interval $[0, T]$

$$\frac{Y^{(c)}}{c} \Rightarrow Y^{(\infty)} \quad \text{as } c \rightarrow \infty.$$

PROOF OF THEOREM 5.1. We will prove this result using a coupling argument. In between document arrival-departure events, a MQNwS behaves as a closed queueing network. We couple MQNwS so that in between arrival-departure events this closed queueing network behaviour is determined by a single closed queueing process. By doing this, Skorohod convergence results become a consequence of renewal theory results.

We split the proof into four sections. In the first section, we couple the queueing network’s initial states. In the second section, we state an induction hypothesis which we will use to prove weak convergence. In the third section, we form a coupling of our queueing networks. In the fourth section, we prove this coupling satisfies the induction hypothesis. Finally in the fifth section, we prove weak convergence in the Skorohod topology.

Coupling the initial state:

We start by coupling the initial state of our process. By (18) and the Skorohod Representation Theorem [3, Section 6] we may choose a sequence of coupled random variables $\{Y^{(c)}(0)\}_{c \in \mathbb{N} \cup \{\infty\}}$ such that, almost surely

$$(19) \quad \frac{Y^{(c)}(0)}{c} \xrightarrow{c \rightarrow \infty} Y^{(\infty)}(0).$$

For $c \in \mathbb{N}$ and given our coupled sequence $\{Y^{(c)}(0)\}_{c \in \mathbb{N} \cup \{\infty\}}$ we know the required distribution of $Q^{(c)}(0)$. We may choose a sequence of functions $f^{(c)} : \mathcal{Y} \times [0, 1] \rightarrow \mathcal{Q}$ such that,

$$\mathbb{P}(f^{(c)}(y, U) = q) = \mathbb{P}(Q^{(c)}(0) = q | Y^{(c)}(0) = y), \quad \forall q \in \mathcal{Q}, y \in \mathcal{Y}$$

where here U is an independent uniform random variable on $[0, 1]$. In particular, we specify the function $f^{(c)}(y, \cdot) : [0, 1] \rightarrow \mathcal{Q}$ to be such that $f^{(c)}(y, u) = q$ for values of u belonging to an interval of length $\mathbb{P}(Q^{(c)}(0) = q | Y^{(c)}(0) = y)$. Thus from a single uniform random variable and the coupled sequence $\{Y^{(c)}(0)\}_{c \in \mathbb{N}}$, we may define the coupled initial state of each MQNwS by

$$(20) \quad Q^{(c)}(0) = f^{(c)}(Y^{(c)}(0), U), \quad c \in \mathbb{N}.$$

Induction Hypothesis:

We now inductively construct our coupled process under the following induction hypothesis on $\kappa \in \mathbb{Z}_+$. For $c \in \mathbb{N} \cup \{\infty\}$, let $\tau^{\kappa, (c)}$ be the κ -th document arrival-departure event for the flow level state of our coupled process $Y^{(c)}$, $c \in \mathbb{N} \cup \{\infty\}$. We assume under this induction hypothesis that we have already defined $Y^{(c)}$ on the interval $[0, \tau^{\kappa, (c)}]$ and that under this coupling

$$(21) \quad \tau^{k, (c)} \xrightarrow{c \rightarrow \infty} \tau^{k, (\infty)}, \quad k = 0, \dots, \kappa,$$

$$(22) \quad \frac{Y^{(c)}(\tau^{k, (c)})}{c} \xrightarrow{c \rightarrow \infty} Y^{(\infty)}(\tau^{k, (\infty)}), \quad k = 0, \dots, \kappa.$$

Our induction hypothesis states that there exists a coupling of $Y^{(c)}$ extended to the next arrival-departure event, i.e. on the interval $(\tau^{\kappa, (c)}, \tau^{\kappa+1, (c)} \wedge T]$, such that,

$$(23) \quad \sup_{t \in (\tau^{\kappa, (c)}, \tau^{\kappa+1, (c)} \wedge T]} |\lambda^{\kappa, (c)}(t) - t| \xrightarrow{c \rightarrow \infty} 0$$

$$(24) \quad \sup_{t \in (\tau^{\kappa, (c)}, \tau^{\kappa+1, (c)} \wedge T]} \left\| Y^{(\infty)}(\lambda^{\kappa, (c)}(t)) - \frac{Y^{(c)}(t)}{c} \right\| \xrightarrow{c \rightarrow \infty} 0,$$

where $\lambda^{\kappa, (c)} : [\tau^{\kappa, (c)}, \tau^{\kappa+1, (c)} \wedge T] \rightarrow [\tau^{\kappa, (\infty)}, \tau^{\kappa+1, (\infty)} \wedge T]$ is the function that linearly interpolates between $\lambda^{\kappa, (c)}(\tau^{\kappa, (c)}) = \tau^{\kappa, (\infty)}$ and $\lambda^{\kappa, (c)}(\tau^{\kappa+1, (c)} \wedge T) = \tau^{\kappa+1, (\infty)} \wedge T$. Here norm $\|\cdot\|$ is defined by (5). This completes the statement of the induction hypothesis.

Note taking $\tau^{0, (c)} = 0$, $\forall c \in \mathbb{N}$ by (19) and (20) our induction hypothesis holds for $\kappa = 0$. Also note the convergence statements (23) and (24) are stronger than (21) and (22).

Coupling:

Given our induction hypothesis holds until time $\tau^{\kappa, (c)}$, we will define a coupling until the next arrival-departure time $\tau^{\kappa+1, (c)}$. In order to simplify notation, we will use the shorthand $\tilde{q}^{\kappa, (c)} = \tilde{Q}^{(c)}(\tau^{\kappa, (c)})$, $y^{\kappa, (c)} = Y^{(c)}(\tau^{\kappa, (c)})$

and $n^{\kappa,(c)} = N^{(c)}(\tau^{\kappa,(c)})$. These denote the packet level state, flow level state and number of documents in transfer for the c -th network at time $\tau^{\kappa,(c)}$.

By assumption (21), $\exists c'$ such that $\forall c > c'$

$$(25) \quad n^{\kappa,(c)} = n^{\kappa,(\infty)}.$$

Let \bar{Q}^κ define a closed multi-class queueing network with $n^{\kappa,(\infty)}$ packets across each route and with queue service capacities defined by $(\phi_j(\cdot) : j \in \mathcal{J})$. For states $\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})$, let $\sigma_{\tilde{q}}$ define the first time \bar{Q}^κ hits the state \tilde{q} . As \bar{Q}^κ is an irreducible, positive recurrent Markov chain, almost surely $\sigma_{\tilde{q}} < \infty$, $\forall \tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})$.

The packet level state of a multi-class queueing network with spinning behaves as a closed queueing network between arrival-departure events. Thus we can extend the packet level description of the c -th multi-class queueing network with spinning by defining, $\forall c > c'$

$$(26) \quad \tilde{Q}^{(c)}(t) = \bar{Q}^\kappa(c(t - \tau^{\kappa,(c)}) + \sigma_{\tilde{q}^{\kappa,(c)}}), \quad t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)}).$$

We will shortly define $\tau^{\kappa+1,(c)}$.

We, also, define the flow level state of the c -th multi-class queueing network with spinning. We associate each packet in the closed queueing network \bar{Q}^κ at time $\sigma_{\tilde{q}^{\kappa,(c)}}$ with a packet in the c -th MQNwS at time $\tau^{\kappa,(c)}$. For each route, let k index the packets associated with each document at time $\tau^{\kappa,(c)}$. We retain this same index until time $\tau^{\kappa+1,(c)}$. Let $\bar{A}_{ik}^\kappa(t)$ denote the number of transitions where the k -th packet on route i in the c -th MQNwS has traversed route i in closed queueing network $\bar{Q}^\kappa(t)$ by time t . We define the components of the flow level process of the c -th MQNwS by

$$(27) \quad Y_{ik}^{(c)}(t) = Y_{ik}^{(c)}(\tau^{\kappa,(c)}) - \bar{A}_{ik}^\kappa(c\{t - \tau^{\kappa,(c)}\} + \sigma_{\tilde{q}^{\kappa,(c)}}),$$

for $k = 1, \dots, n_i^{\kappa,(c)}$, $i \in \mathcal{I}$, $t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)})$, and for $c > c'$. Similarly, for $c = \infty$, for the spinning network we define

$$(28) \quad Y_{ik}^{(\infty)}(t) = Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}) - \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}}(t - \tau^{\kappa,(\infty)}),$$

for $k = 1, \dots, n_i^{\kappa,(\infty)}$, $i \in \mathcal{I}$ and $t \in (\tau^{\kappa,(\infty)}, \tau^{\kappa+1,(\infty)})$. Recall in the definition of the flow level state of a MQNwS, residual file sizes of each route are indexed to be increasing in size, (i.e. $y_{ik} \leq y_{i,k+1}$). In both expressions (27) and (28) we do not index $Y_{ik}^{(c)}$ so that residual file sizes are increasing. Instead we index $Y_{ik}^{(c)}$ so that it is associated with a specific packet on route i in

closed queueing network \bar{Q}^κ . This representation is required so that packet indices do not change over interval $(\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)})$. Even so, these indices are a permutation of the ordering in which residual files sizes are increasing in size.

Note if the processing of the (i, k) -th document is not interrupted by another arrival departure event then, for $c > c'$, this document would depart at time

$$S_{ik}^{\kappa,(c)} = \tau^{\kappa,(c)} + \inf\{t : \bar{A}_{ik}^{\kappa,(c)}(ct + \sigma_{\bar{q}^{\kappa,(c)}}) = Y_{ik}^{(c)}(\tau^{\kappa,(c)})\},$$

and, for $c = \infty$, this would occur at time

$$(29) \quad S_{ik}^{\kappa,(\infty)} = \tau^{\kappa,(\infty)} + \frac{n_i^{\kappa,(\infty)}}{\Lambda_i^{SN}(n^{\kappa,(\infty)})} Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}).$$

In addition, for each $i \in \mathcal{I}$, let E_i^κ be an independent exponential random variable with mean ν_i^{-1} . We define $E_i^{\kappa,(c)}$ by

$$E_i^{\kappa,(c)} = \tau^{\kappa,(c)} + E_i^\kappa, \quad c \in \mathbb{N} \cup \{\infty\}.$$

$E_i^{\kappa,(c)}$ denotes the next arrival of a i document assuming it is uninterrupted by another arrival departure event. From these terms, we can define the next arrival-departure event by

$$(30) \quad \tau^{\kappa+1,(c)} := \min \left(\{S_{ik}^{\kappa,(c)} : k = 1, \dots, N_i^{(c)}, i \in \mathcal{I}\} \cup \{E_i^{\kappa,(c)} : i \in \mathcal{I}\} \right), \\ c \in \mathbb{N} \cup \{\infty\}$$

which arrival-departure event occurs depends on which term minimises this term. Note by Assumption 4, for each $c \in \mathbb{N} \cup \{\infty\}$ there is always a unique minimum of this term.

By the packet level coupling (26) and the flow level coupling (27-28), we have coupled our processes on the interval $(\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)})$. We now include the transition at time $\tau^{\kappa+1,(c)}$: if (30) is minimised by $S_{ik}^{\kappa,(c)}$, we define the c -th multi-class queueing network at time $\tau^{\kappa+1,(c)}$ by appropriately removing the (i, k) document and packet from the network's state description at time $\tau^{\kappa+1,(c)-}$, and if (30) is minimised by $E_i^{\kappa,(c)}$ then we add a new document and packet to the flow state and packet state of the system, this document will be of (residual) size $X_i^{\kappa+1,(c)} \sim X_i^{(c)}$. Here $\{X_i^{\kappa+1,(c)}\}_{c \in \mathbb{N} \cup \{\infty\}}$ is an independent sequence of random variables satisfying,

$$(31) \quad \frac{X_i^{\kappa+1,(c)}}{c} \xrightarrow{c \rightarrow \infty} X_i^{\kappa+1,(\infty)}.$$

This completes the coupling of our process on the interval $(\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)})$.

Proof of induction step:

Given our coupling up to time $\tau^{\kappa+1,(c)}$, we now concern ourselves with proving the convergence statements (23) and (24). The following three lemmas will help to verify this.

LEMMA A. *Almost surely, for $i \in \mathcal{I}$, $\eta > 0$, $k = 1, \dots, n_i$*

$$\sup_{t \in [0, \eta]} \left| \frac{\bar{A}_{ik}^{\kappa}(ct + \sigma_{\tilde{q}^{\kappa,(c)}})}{c} - \frac{\Lambda_i^{SN}(n^{\kappa,(c)})}{n_i^{\kappa,(c)}} t \right| \xrightarrow{c \rightarrow \infty} 0.$$

PROOF OF LEMMA A. We consider c sufficiently large so that (25) holds. Let $R_{ik,\tilde{q}}(t)$ be the number times by time t the k -th route i packet has completed its route in the closed queueing network \bar{Q}^{κ} , when the closed queueing network was in state $\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})$. Let $\gamma_{ik}(\tilde{q})$ be the drift of $R_{ik,\tilde{q}}$. For any Markov chain, the process that records the current state of the Markov chain and the next state, is also a Markov chain. So $R_{ik,\tilde{q}}$ is a renewal process and thus obeys the Functional Renewal Theorem. This gives that, almost surely, for all $\eta > 0$

$$\max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i^{\kappa,(\infty)}} \max_{\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})} \sup_{t \in [0, \eta]} \left| \frac{R_{ik,\tilde{q}}(ct)}{c} - \gamma_{ik}(\tilde{q})t \right| \xrightarrow{c \rightarrow \infty} 0.$$

For a proof of the Functional Renewal Theorem, see [5, page 106]. By the definition of $\bar{A}_{ik}^{\kappa}(t)$ and Corollary 3.5, we know that,

$$\bar{A}_{ik}^{\kappa}(t) = \sum_{\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})} R_{ik,\tilde{q}}(t) \quad \text{and} \quad \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} = \sum_{\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})} \gamma_{ik}(\tilde{q}).$$

So, noting that $\tilde{Q}(n^{\kappa,(\infty)})$ is a finite set, we have that, almost surely, $\forall \eta > 0$

$$\begin{aligned} & \sup_{t \in [0, \eta]} \left| \frac{\bar{A}_{ik}^{\kappa,(c)}(ct)}{c} - \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} t \right| \\ & \leq \sum_{\tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})} \max_{r \in \mathcal{I}} \max_{k'=1, \dots, n_r^{\kappa,(\infty)}} \sup_{t \in [0, \eta]} \left| \frac{R_{rk',\tilde{q}}(ct)}{c} - \gamma_{rk'}(\tilde{q})t \right| \xrightarrow{c \rightarrow \infty} 0. \end{aligned}$$

As \bar{Q}^{κ} is recurrent on all states in $\tilde{Q}(n^{\kappa,(\infty)})$, almost surely, $\sigma_{\tilde{q}} < \infty \forall \tilde{q} \in \tilde{Q}(n^{\kappa,(\infty)})$. Thus, almost surely,

$$\sup_{t \in [0, \eta]} \left| \frac{\bar{A}_{ik}^{\kappa}(ct + \sigma_{\tilde{q}^{\kappa,(c)}})}{c} - \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} \left(t + \frac{\sigma_{\tilde{q}^{\kappa,(c)}}}{c} - \frac{\sigma_{\tilde{q}^{\kappa,(c)}}}{c} \right) \right|$$

$$\leq \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)}) \sigma_{\bar{q}^{\kappa,(c)}}}{n_i^{\kappa,(\infty)} c} + \sup_{t \in [0, \eta + \sigma_{\bar{q}^{\kappa,(c)}}]} \left| \frac{\bar{A}_{ik}^\kappa(ct)}{c} - \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} t \right| \xrightarrow{c \rightarrow \infty} 0.$$

This completes the proof of Lemma A. \square

This renewal result characterises the limiting behaviour of $S_{ik}^{\kappa,(c)}$ (the time until document k 's departure given the current flow level state).

LEMMA B. *For each $i \in \mathcal{I}$ and $k = 1, \dots, n_i^{\kappa,(\infty)}$, almost surely*

$$S_{ik}^{\kappa,(c)} \xrightarrow{c \rightarrow \infty} S_{ik}^{\kappa,(\infty)}.$$

PROOF OF LEMMA B. By Lemma A and induction hypothesis (22), almost surely, $\forall \epsilon > 0$ and $\forall \eta > 0$ such that $\eta > S_{ik}^{\kappa,(\infty)} - \tau^{\kappa,(c)} + \frac{2\epsilon n_i^{\kappa,(\infty)}}{\Lambda_i^{SN}(n^{\kappa,(\infty)})}$, $\exists c'$ such that $\forall c > c'$,

$$(32) \quad \sup_{t \in [0, \eta]} \left| \frac{\bar{A}_{ik}^{\kappa,(\infty)}(ct + \sigma_{\bar{q}^{\kappa,(c)}}) - \frac{\Lambda_i^{SN}(n^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} t}{c} \right| < \epsilon,$$

$$(33) \quad \left| \frac{Y_{ik}^{(c)}(\tau^{\kappa,(c)})}{c} - Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}) \right| < \epsilon.$$

Hence, recalling the definition of $S_{ik}^{\kappa,(\infty)}$ in (29), the above two inequalities imply for all documents (i, k)

$$\begin{aligned} & \frac{1}{c} \bar{A}_{ik}^\kappa \left(c \left\{ S_{ik}^{\kappa,(\infty)} - \tau^{\kappa,(\infty)} - \frac{2\epsilon n_i^{\kappa,(\infty)}}{\Lambda_i^{SN}(n^{\kappa,(\infty)})} \right\} + \sigma_{\bar{q}^{\kappa,(c)}} \right) \\ & \leq Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}) - \epsilon < \frac{Y_{ik}^{(c)}(\tau^{\kappa,(c)})}{c}. \end{aligned}$$

Thus

$$\begin{aligned} S_{ik}^{(c)} - \tau^{\kappa,(c)} &= \inf\{t \geq 0 : \bar{A}_{ik}^\kappa(ct + \sigma_{\bar{q}^{\kappa,(c)}}) = Y_{ik}^{(c)}(\tau^{\kappa,(c)})\} \\ &> S_{ik}^{(\infty)} - \tau^{\kappa,(\infty)} - \frac{2\epsilon n_i^{\kappa,(\infty)}}{\Lambda_i^{SN}(n^{\kappa,(\infty)})}. \end{aligned}$$

By a similar argument one can see that

$$S_{ik}^{(c)} - \tau^{\kappa,(c)} < S_{ik}^{(\infty)} - \tau^{\kappa,(\infty)} + \frac{2\epsilon n_i^{\kappa,(\infty)}}{\Lambda_i^{SN}(n^{\kappa,(\infty)})}.$$

Thus $S_{ik}^{(c)} - \tau^{\kappa,(c)} \rightarrow S_{ik}^{(\infty)} - \tau^{\kappa,(\infty)}$ as $c \rightarrow \infty$, almost surely. Thus as we assume (21) holds, almost surely

$$S_{ik}^{(c)} \xrightarrow{c \rightarrow \infty} S_{ik}^{(\infty)}.$$

This completes the proof of Lemma B. \square

Recall, that we had not ordered elements $Y_{ik}^{(c)}$ in increasing order, instead we indexed $Y_{ik}^{(c)}$ to be associated with each individual packet being processed in the closed queueing network \bar{Q}^κ . The following lemma helps us re-associate the desired increasing ordering of the terms $Y_{ik}^{(c)}$.

LEMMA C. *Let $y, y' \in \mathbb{R}_+^n$ be such that $y_1 \leq \dots \leq y_n$, $y'_1 \leq \dots \leq y'_n$ and let $p : \{1, \dots, n\} \rightarrow \{1, \dots, n\}$ be a permutation, then*

$$(34) \quad \max_{k=1, \dots, n} |y_k - y'_k| \leq \max_{k=1, \dots, n} |y_k - y'_{p(k)}|.$$

PROOF OF LEMMA C. We prove the result by induction on n , under the induction hypothesis that for all $\delta > 0$

$$(35) \quad |y_k - y'_{p(k)}| < \delta \quad \forall k = 1, \dots, n \quad \text{implies} \quad |y_k - y'_k| < \delta \quad \forall k = 1, \dots, n.$$

The hypothesis clearly holds for $n = 1$. Suppose the induction hypothesis holds for $n - 1$. Take $i = p^{-1}(n)$ and $j = p(n)$. If $i = j (= n)$ then the problem clearly reduces to the $n - 1$ case. Assume $i \neq j$. We know

$$y_n \geq y_i \quad y'_n \geq y'_j.$$

Also, if

$$|y_n - y'_j| < \delta \quad \text{and} \quad |y_i - y'_n| < \delta$$

then

$$(36) \quad \begin{aligned} y_n \geq y_i > y'_n - \delta \quad \text{and} \quad y'_n \geq y'_j > y_n - \delta \\ \text{therefore} \quad |y_n - y'_n| < \delta. \end{aligned}$$

Similarly,

$$(37) \quad \begin{aligned} y_i > y'_n - \delta \geq y'_k \quad \text{and} \quad y'_j > y_n - \delta \geq y_i - \delta \\ \text{therefore} \quad |y_i - y'_j| < \delta. \end{aligned}$$

We can now define a new permutation on $\{1, \dots, n-1\}$,

$$p'(k) = \begin{cases} p(k) & \text{if } k \neq i \\ j & \text{if } k = i. \end{cases}$$

Since (37) holds we have reduced this problem to a problem on $n-1$ variables with equality (36) still holding. This completes the proof of our induction hypothesis. Since δ is arbitrary it is clear that (35) is equivalent to (34). This completes the proof of Lemma C. \square

By Lemma B, induction hypothesis (21) and the definition of $\tau^{k+1,(c)}$ (27-28), we know that

$$\tau^{\kappa+1,(c)} \rightarrow \tau^{\kappa+1,(\infty)} \quad \text{as } c \rightarrow \infty.$$

By Assumption 4, $\tau^{\kappa+1,(\infty)}$ is achieved by a distinct minimum and consequently there exists a c'' such that $\forall c > c''$ the argument attaining $\tau^{\kappa+1,(c)}$ in (30) is the same as that attaining $\tau^{\kappa+1,(\infty)}$. Thus the coupled processes will have the same document arrival-departure event occur at time $\tau^{\kappa+1,(c)}$, $\forall c \in \{c'' + 1, \dots, \infty\}$.

We can now verify (23) from the induction hypothesis:

$$(38) \quad \begin{aligned} & \sup_{t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)} \wedge T]} |\lambda^{\kappa,(c)}(t) - t| \\ & = |\tau^{\kappa,(\infty)} - \tau^{\kappa,(c)}| \vee |\tau^{\kappa+1,(\infty)} \wedge T - \tau^{\kappa+1,(c)} \wedge T| \xrightarrow{c \rightarrow \infty} 0. \end{aligned}$$

We can also prove (24) from the induction hypothesis. We use the following set of inequalities which will subsequently be explained, $\forall c > c''$,

(39)

$$\sup_{t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)} \wedge T]} \left\| Y^{(\infty)}(\lambda^{\kappa,(c)}(t)) - \frac{Y^{(c)}(t)}{c} \right\|$$

(40)

$$\begin{aligned} & \leq \max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i^{\kappa,(\infty)}} \sup_{t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)} \wedge T]} \left| Y_{ik}^{(\infty)}(\lambda^{\kappa,(c)}(t)) - \frac{Y_{ik}^{(c)}(t)}{c} \right| \\ & \quad + \max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i^{\kappa,(\infty)}} \left| X_{ik}^{\kappa+1,(\infty)} - \frac{X_{ik}^{\kappa+1,(c)}}{c} \right| \\ & \leq \max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i^{\kappa,(\infty)}} \left[\left| Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}) - \frac{Y_{ik}^{(c)}(\tau^{\kappa,(c)})}{c} \right| \right] \end{aligned}$$

$$\begin{aligned}
 & + \frac{\Lambda_i^{SN}(n_i^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} \times \sup_{t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)} \wedge T]} |\lambda^{\kappa,(c)}(t) - t| \\
 & + \sup_{t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)})} \left| \frac{\Lambda_i^{SN}(n_i^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} (t - \tau^{\kappa,(c)}) - \frac{\bar{A}_{ik}^\kappa(c\{t - \tau^{\kappa,(c)}\} + \sigma_{\bar{q}^{\kappa,(c)}})}{c} \right| \\
 & \quad \left. + \frac{\Lambda_i^{SN}(n_i^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} |\tau^{\kappa,(c)} - \tau^{\kappa,(\infty)}| \right] \\
 (41) \quad & + \max_{i \in \mathcal{I}} \max_{k=1, \dots, n_i^{\kappa,(\infty)}} \left| X_{ik}^{\kappa+1,(\infty)} - \frac{X_{ik}^{\kappa+1,(c)}}{c} \right| \\
 & \xrightarrow{c \rightarrow \infty} 0
 \end{aligned}$$

In the first inequality, we apply Lemma C so that we index each packet according to its position within closed queueing network \bar{Q}^κ (see description of (27) and (26)). This is so residual file sizes $Y_{ik}^{(c)}(t)$ are not necessarily indexed to be increasing. Also the first inequality over estimates (39) by including the file sizes of all possible arrivals that could occur at time $\tau^{\kappa+1,(c)}$. In the second inequality we apply the triangle inequality to (40) by using the two facts,

$$\begin{aligned}
 Y_{ik}^{(\infty)}(\lambda^{\kappa,(c)}(t)) &= Y_{ik}^{(\infty)}(\tau^{\kappa,(\infty)}) \\
 &\quad - \frac{\Lambda_i^{SN}(n_i^{\kappa,(\infty)})}{n_i^{\kappa,(\infty)}} \left(\{\lambda^{\kappa,(c)}(t) - t\} + \{t - \tau^{\kappa,(c)}\} + \{\tau^{\kappa,(c)} - \tau^{\kappa,(\infty)}\} \right) \\
 Y_{ik}^{(c)}(t) &= Y_{ik}^{(c)}(\tau^{\kappa,(c)}) - \bar{A}_{ik}^\kappa(c\{t - \tau^{\kappa,(c)}\} + \sigma_{\bar{q}^{\kappa,(c)}}).
 \end{aligned}$$

for $t \in (\tau^{\kappa,(c)}, \tau^{\kappa+1,(c)} \wedge T]$. The first expression in equation (41) converges to 0 by induction assumption (22); the second expression converges to 0 by (38); the third term converges by (21); the fourth converges by Lemma A and fifth term converges by expression (31). We have thus demonstrated (23) and (24) hold. This verifies our induction hypothesis.

Our induction argument is sufficient to couple our process on interval $[0, T]$. Since there are almost surely a finite number of documents in transfer at time $t = 0$ and a finite number of document arrivals in interval $[0, T]$, it must be that

$$\{\kappa : \tau^{\kappa,(\infty)} < T\} \text{ is bounded almost surely.}$$

Since we have proven that almost surely $\tau^{\kappa,(c)} \rightarrow \tau^{\kappa,(\infty)}$ as $c \rightarrow \infty$, for all $\kappa \in \mathbb{Z}_+$,

$$\{\kappa : \tau^{\kappa,(c)} < T\} \text{ is uniformly bounded over } c \in \mathbb{N} \cup \{\infty\} \text{ almost surely.}$$

Thus by our inductive argument, we may couple our process $\{Y^{(c)}\}_{c \in \mathbb{N} \cup \{\infty\}}$ on the interval $[0, T]$.

Skorohod convergence:

Taking $\lambda^{(c)}(t) = \lambda^{\kappa, (c)}(t)$ for $t \in [\tau^{\kappa, (c)}, \tau^{\kappa+1, (c)} \wedge T]$. We have by statements (23) and (24) that

$$\begin{aligned} \sup_{t \in [0, T]} |\lambda^{(c)}(t) - t| &\xrightarrow{c \rightarrow \infty} 0 \\ \sup_{t \in [0, T]} \left\| Y^{(\infty)}(\lambda^{(c)}(t)) - \frac{Y^{(c)}(t)}{c} \right\| &\xrightarrow{c \rightarrow \infty} 0. \end{aligned}$$

Thus, almost surely, we have convergence in the Skorohod topology on $[0, T]$,

$$\frac{Y^{(c)}}{c} \xrightarrow{c \rightarrow \infty} Y^{(\infty)}.$$

Since, the Skorohod convergence occurs almost surely in this coupling, for all continuous bounded functions $f : D[0, t] \rightarrow \mathbb{R}_+$,

$$f\left(\frac{Y^{(c)}}{c}\right) \xrightarrow{c \rightarrow \infty} f(Y^{(\infty)}), \quad \text{almost surely.}$$

Thus by the Bounded Convergence Theorem,

$$\mathbb{E}f\left(\frac{Y^{(c)}}{c}\right) \xrightarrow{c \rightarrow \infty} \mathbb{E}f(Y^{(\infty)})$$

or, in other words, $\frac{Y^{(c)}}{c}$ converges weakly to $Y^{(\infty)}$ in the Skorohod topology. This completes the proof of Theorem 5.1. \square

6. Insensitivity of spinning networks. The insensitivity of the spinning network is a consequence of Theorem 5.1. To prove this we will first require two technical lemmas.

LEMMA 6.1. *Let $(X^{(c)} : c \in \mathbb{N} \cup \{\infty\})$ be a sequence of random variables. Let $X^{(c)}$ have values in \mathbb{N} and mean $\frac{c}{\mu}$ for $c \in \mathbb{N}$. Let $X^{(\infty)}$ have values in \mathbb{R}_+ and mean $\frac{1}{\mu}$. Define random variables $(\bar{X}^{(c)} : c \in \mathbb{N} \cup \{\infty\})$ by*

$$\begin{aligned} \mathbb{P}(\bar{X}^{(c)} \leq y) &= \frac{\mu}{c} \sum_{z=1}^y \mathbb{P}(X^{(c)} \geq z), \quad c \in \mathbb{N}, \\ (42) \quad \mathbb{P}(\bar{X}^{(\infty)} \leq y) &= \mu \int_0^y \mathbb{P}(X^{(\infty)} \geq z) dz. \end{aligned}$$

If

$$\frac{X^{(c)}}{c} \Rightarrow X^{(\infty)} \quad \text{as } c \rightarrow \infty,$$

then

$$(43) \quad \frac{\bar{X}^{(c)}}{c} \Rightarrow \bar{X}^{(\infty)} \quad \text{as } c \rightarrow \infty.$$

PROOF. $\mathbb{P}(X^{(\infty)} \geq z)$ can only have countably many points of discontinuity. Thus by integration by substitution and the Bounded Convergence Theorem, we have that, for all $y \in \mathbb{R}_+$,

$$\begin{aligned} \mathbb{P}\left(\frac{\bar{X}^{(c)}}{c} \leq y\right) &= \mathbb{P}\left(\bar{X}^{(c)} \leq \lfloor cy \rfloor\right) = \frac{\mu}{c} \int_1^{\lfloor cy \rfloor} \mathbb{P}(X^{(c)} \geq z) dz \\ &= \mu \int_{1/c}^{\lfloor cy \rfloor / c} \mathbb{P}\left(\frac{X^{(c)}}{c} \geq z\right) dz \xrightarrow{c \rightarrow \infty} \mu \int_0^y \mathbb{P}(X^{(\infty)} \geq z) dz = \mathbb{P}(\bar{X}^{(\infty)} \leq y). \end{aligned}$$

□

To prove Theorem 5.1, we assumed no simultaneous arrival-departure events occurred. We now demonstrate that these assumptions hold for the case of a spinning network, with non-atomic document sizes.

LEMMA 6.2. *Suppose the initial distribution $Y(0)$ conditional of $N(0)$ consists of independent non-atomic random variables Y_{i_k} $k = 1, \dots, N_i(0), i \in \mathcal{I}$. Given documents size distributions $X_i, i \in \mathcal{I}$ are non-atomic, then, almost surely,*

- a) *There are no simultaneous document arrival-departure events, i.e. Assumption 4 holds.*
- b) *For all $t \in \mathbb{R}_+$, almost surely, no document arrival-departure event occurs at time t .*

PROOF. Let $Y = (Y_t : t \in \mathbb{R}_+)$ be the spinning networks flow level process description. Since arrivals A_1, A_2, \dots form a Poisson process almost surely no two arrivals occur at the same time and for each $t \in \mathbb{R}_+$ almost surely no arrival occurs at time t . Since exponential random variable $A_k - A_{k-1}$ is independent of $(Y_t : t \leq A_{k-1})$, there is zero probability that an arrival A_k coincides with departures. Therefore, an arrival cannot coincide with a departure.

It remains to show that no two document departures may occur simultaneously. Let D_k be the departure of some document k of initial size X_k (or

initial residual size Y_{ik} at time 0). Let A_k be that document's arrival time (take $A_k = 0$ if the document is present at time zero). Let Y' be the process derived from Y in which document k never departs the BSN (i.e. behaving as if $X_k = \infty$). Note that $Y'(t) = Y(t)$ for all $t < D_k$ and D_k coincides with a document departure in Y iff D_k coincides with a departure in Y' . Note that as that X_k is conditionally independent of $(Y'(t) : t > A_k)$ conditional on $(Y'(t) : t \leq A_k)$ and D_k is non-atomic as it is a strictly increasing function of non-atomic independent random variable X_k . Thus, conditional on $(Y'(t) : t \leq A_k)$ the probability that non-atomic random variable D_k coincides with the countable set of departure events in $(Y'(t) : t > A_k)$ or at a specific time $t \in \mathbb{R}_+$ is zero. Thus, the probability two departure events coincide is zero and the probability that departure occurs at a specific time t is zero. \square

We can now prove one of the main results of this chapter: the insensitivity of the spinning network.

COROLLARY 6.1. *Given Assumption 1, the spinning network has a stationary distribution which is insensitive to all non-atomic document size distributions.*

PROOF. We can take document sizes $X_i^{(c)}$ such that $\mathbb{E}X_i^{(c)} = \frac{c}{\mu_i}$ and $\frac{X_i^{(c)}}{c} \Rightarrow X_i^{(\infty)}$. As in Theorem 5.1, we consider a sequence of multi-class queueing networks with spinning associated with these document size distributions and with queue service rates $c\phi_j(\cdot)$. From Corollary 3.2 and Corollary 3.3, the prelimit stationary distribution of $Y^{(c)}$ and $N^{(c)}$, $c \in \mathbb{N}$ is

$$(44) \quad \mathbb{P}(Y^{(c)}(0) = y) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \binom{n_i}{n_{iy} : y \in \mathbb{N}} \prod_{k=1}^{n_i} (\nu_i \mathbb{P}(X_i \geq y_{ik})), \quad \forall y \in \mathcal{Y}, t \in \mathbb{R}_+,$$

$$(45) \quad \mathbb{P}(N^{(c)}(0) = n) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \rho_i^{n_i}, \quad \forall n \in \mathbb{Z}_+^I.$$

We can construct $Y^{(c)}(0)$, by taking a vector $N(0)$ according to distribution (45) then, for each $i \in \mathcal{I}$ and $k = 1, \dots, N_i(0)$, $Y_{ik}^{(c)}(0)$ is taken by selecting and ordering independent random variables $\bar{X}_i^{(c)}$, where $\bar{X}_i^{(c)}$ is defined from $X_i^{(c)}$ by (42). Given Lemma 6.1,

$$(46) \quad \frac{Y^{(c)}(0)}{c} \Rightarrow Y^{(\infty)}(0) \quad \text{as} \quad c \rightarrow \infty,$$

where $Y^{(\infty)}(0)$ has density

$$(47) \quad \frac{B_n}{B} \prod_{i \in \mathcal{I}} n_i! \prod_{k=1}^{n_i} (\rho_i \mathbb{P}(\bar{X}_i^{(\infty)} \in dx_{ik})),$$

and also

$$(48) \quad bP(N^{(\infty)}(0) = n) = \frac{B_n}{B} \prod_{i \in \mathcal{I}} \rho_i^{n_i},$$

$\forall n \in \mathbb{Z}_+^I$ and for $x_{ik} \leq x_{ik+1}$ $k = 1, \dots, n_i - 1$, $i \in \mathcal{I}$. Note the above expression for $N^{(\infty)}$ depends on $X_i^{(\infty)}$ only through its mean. Hence if (47) is the stationary distribution for our limit process then this distribution must be insensitive.

We now show that (47) provides a stationary distribution. It is known that if a sequence of processes weakly converge in the Skorohod topology and if, almost surely, there is not jump at time t then the marginal distribution at time t must weakly converge, see [3, Theorem 12.5]. By Lemma 6.2, almost surely no jump occurs at time t for $N^{(\infty)}$ and, by Theorem 5.1, $N^{(c)} \Rightarrow N^{(\infty)}$ as $c \rightarrow \infty$ in the Skorohod topology. Thus, $N^{(c)}(t) \Rightarrow N^{(\infty)}(t)$ i.e. the marginal distributions converge at time t . Thus, when processes $Y^{(c)}$, $c \in \mathbb{N}$ are stationary, by for any continuous bounded function $f : \mathcal{Y} \rightarrow \mathbb{R}$

$$\mathbb{E}f(Y^{(\infty)}(0)) = \lim_{c \rightarrow \infty} \mathbb{E}f\left(\frac{Y^{(c)}(0)}{c}\right) = \lim_{c \rightarrow \infty} \mathbb{E}f\left(\frac{Y^{(c)}(t)}{c}\right) = \mathbb{E}f(Y^{(\infty)}(t)).$$

The first equality holds by (46); the second holds by the stationarity of $Y^{(c)}$; and the third holds by the weak convergence of the marginal distributions. This proves (47) gives a stationary distribution of the spinning network, and consequently, from (45) and (48), we see that the spinning network is insensitive. \square

REFERENCES

- [1] BARBOUR, A. D. (1976). Networks of queues and the method of stages. *Advances in Applied Probability*, **8** pp. 584–591. URL <http://www.jstor.org/stable/1426145>. MR0440729
- [2] BASKETT, F., CHANDY, K. M., MUNTZ, R. R. and PALACIOS, F. G. (1975). Open, closed, and mixed networks of queues with different classes of customers. *J. ACM*, **22** 248–260. MR0365749
- [3] BILLINGSLEY, P. (1999). *Convergence of Probability Measures*. Wiley, New York. MR1700749
- [4] BONALD, T. and PROUTIERE, A. (2004). On performance bounds for balanced fairness. *Performance Evaluation*, **55** 25–50.

- [5] CHEN, H. and YAO, D. D. (2001). *Fundamentals of Queueing Networks: Performance, Asymptotics and Optimization*. Springer, New York. [MR1835969](#)
- [6] FEUILLET, M. and ROBERT, P. (2012). A Scaling Analysis of a Transient Stochastic Network (I). *ArXiv e-prints*. [1203.6848](#). [MR2977408](#)
- [7] GROMOLL, H. and WILLIAMS, R. (2009). Fluid limits for networks with bandwidth sharing and general document size distributions. *Annals of Applied Probability*, **19** 243–280. [MR2498678](#)
- [8] KANG, W. N., KELLY, F. P., LEE, N. H. and WILLIAMS, R. J. (2009). State space collapse and diffusion approximation for a network operating under a fair bandwidth sharing policy. *Ann. Appl. Probab.*, **19** 1719–1780. [MR2569806](#)
- [9] KELLY, F. P. (1975). Networks of queues with customers of different types. *Journal of Applied Probability*, **12** 542–554. URL <http://www.jstor.org/stable/3212869>. [MR0388571](#)
- [10] KELLY, F. P. (1979). *Reversibility and Stochastic Networks*. Wiley, Chicester. [MR0554920](#)
- [11] KELLY, F. P. and WILLIAMS, R. J. (2004). Fluid model for a network operating under a fair bandwidth-sharing policy. *Ann. Appl. Probab.*, **14** 1055–1083. [MR2071416](#)
- [12] KLEINROCK, L. (1967). Time-shared systems: a theoretical treatment. *J. ACM*, **14** 242–261. [MR0273994](#)
- [13] M. FROLKOVA, J. R. and ZWART, B. (2012). Limit theorems for bandwidth sharing networks with rate constraints. *Preprint*.
- [14] MASSOULIÉ, L. and ROBERTS, J. (1998). Bandwidth sharing and admission control for elastic traffic. *Telecommunication Systems*, **15** 185–201. [MR1743809](#)
- [15] MASSOULIÉ, L. and ROBERTS, J. (1999). Bandwidth sharing: Objectives and algorithms. *IEEE Infocom 1999*, **10** 320–328.
- [16] PROUTIERE, A. (2003). *Insensitivity and stochastic bounds in queueing networks-Application to flow level traffic modelling in telecommunication networks*. Ph.D. thesis, Ecole Doctorale de l'Ecole Polytechnique.
- [17] VAN DE VEN, P., BORST, S., VAN LEEUWAARDEN, J. and PROUTIERE, A. (2010). Insensitivity and stability of random-access networks. *Performance Evaluation*, **67** 1230–1242. URL <http://www.sciencedirect.com/science/article/pii/S0166531610001082>.
- [18] VIRTAMO, J. (2003). Insensitivity of a network of symmetric queues with balanced service rates. *Internal Report*.
- [19] WALTON, N. S. (2009). Proportional fairness and its relationship with multi-class queueing networks. *Ann. Appl. Probab.*, **22** 2301–2333. [MR2588246](#)
- [20] ZACHARY, S. (2007). A note on insensitivity in stochastic networks. *J. Appl. Probab.*, **44** 238–248. [MR2312999](#)

KORTEWED-DE VRIES INSTITUTE FOR MATHEMATICS
P.O. BOX 94248,
1090 GE AMSTERDAM,
THE NETHERLANDS,
E-MAIL: n.s.walton@uva.nl