

A NOTE ON INSENSITIVITY IN STOCHASTIC NETWORKS

STAN ZACHARY,* *Heriot-Watt University*

Abstract

We give a simple and direct treatment of insensitivity in stochastic networks which is quite general and provides probabilistic insight into the phenomenon. In the case of multi-class networks, our results generalise those of Bonald and Proutière (2002), (2003).

Keywords: Insensitivity; stochastic network; partial balance

2000 Mathematics Subject Classification: Primary 60K20

1. Introduction

It is well known that many stochastic networks – notably queueing and loss networks – have stationary distributions of their level of occupancy which depend on certain input distributions only through the means of the latter. This phenomenon of *insensitivity* has been studied by various authors over an extended period of time, in varying degrees of generality and abstraction, and using a variety of techniques.

In the present paper we revisit this topic to develop an insight of Pechinkin (1983), (1987) which gives a very simple and direct treatment of insensitivity. In particular, our approach avoids those based on brute-force calculations, the consideration of phase-type distributions (see Schassberger (1978), Whittle (1985), and Bonald and Proutière (2002), (2003)), or the use of quite complex machinery for handling generalised semi-Markov processes (see Burman (1981) and Schassberger (1986)) – although such processes are implicit in the current approach. It further avoids assumptions about, for example, continuity of distributions, which are necessary for some of the above approaches, and also explicitly identifies the entire stationary distributions of the networks concerned, showing that, where insensitivity obtains, these stationary distributions have a particularly simple and natural form. Pechinkin used his insight, which involved what is in effect a coupling argument together with induction, to give probabilistic proofs of the insensitivity of a number of single-class loss systems with state-dependent arrival rates – results originally proved analytically by Sevast'yanov (1957). He also indicated the wider applicability of the approach in the single-class case. In the present paper we give a substantial reformulation of the underlying idea, under more general conditions and showing that its most natural expression is in terms of balance equations. This considerably simplifies its application to single-class systems; notably, the quite complex coupling constructions are no longer needed. It further makes possible the extension of the idea to the multi-class networks considered in Section 3. The main aim is to provide probabilistic insight, notably for multi-class networks. Indeed it is shown that insensitivity is simply a by-product, under appropriate conditions, of probabilistic independence.

Received 15 September 2006; revision received 14 November 2006.

* Postal address: Maxwell Institute for Mathematical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK.
Email address: s.zachary@hw.ac.uk

We study networks in which individuals arrive at various *classes* at rates which may depend on the state of the entire system, bringing *workloads* which are independent and identically distributed within classes and which have finite means. Within each class workloads are reduced at rates which may again be state-dependent (when the rate is constant workloads may be identified with lifetimes in classes), and on completion of its workload an individual moves to a different class or leaves the system, with probabilities which may yet again be state-dependent.

In order to obtain insensitivity we typically require that an individual joining a class is immediately served, i.e. has its workload reduced, at a rate which is the same as that of an individual immediately prior to leaving the class (where in each of these cases the number of individuals in each class of the system is the same); more generally, the service discipline should define a network which is *symmetric* in the sense of Kelly (1979). The most common example is that of processor-sharing networks, but other possibilities are well known, for example ‘last-in–first-out preemptive-resume’ networks. We shall concentrate on a very broad class of processor-sharing networks, introduced by Bonald and Proutière (2002) and including, for example, traditional loss networks and processor-sharing Whittle and Jackson networks as special cases. We shall also indicate the simple modifications required to deal with other possibilities.

For the above class of processor-sharing networks, Bonald and Proutière (2002) used phase-type arguments to show that, under conditions which correspond to the satisfaction of the appropriate partial balance equations, the stationary distribution of the number of individuals in each class is insensitive to the workload distributions, subject to the means of the latter being fixed and to the distributions themselves being drawn from the broad class of Cox-type distributions (dense in the class of all distributions on \mathbb{R}_+). In the present paper we formally consider all workload distributions on \mathbb{R}_+ with finite means, and also identify the stationary residual workload distributions. However, as stated above, our main aim is to give a direct and probabilistically natural treatment. It turns out (and is in many cases well known) that, when the appropriate partial balance equations are satisfied for such a network, the stationary distribution of the entire system, *including the specification of residual workloads*, is such that departures from each class are exactly balanced by arrivals to that class (in a sense to be made precise below). Indeed, for single-class systems, this is the essence of Pechinkin’s insight. What is of interest is that the same idea can be extended to establish insensitivity for the very much more general networks considered here, and indeed appears also to establish insensitivity in more abstract settings such as that considered by Whittle (1985), though we do not formally consider this more abstract environment here.

In order to fix ideas, it is convenient to consider first, in Section 2, single-class networks. Here the extension of previous ideas is not too difficult. Nevertheless, it is desirable to give a careful treatment of this case, avoiding notational complexity while preserving rigour, so as both to establish the underlying principle and also to set the scene for the multi-class networks which we consider in Section 3.

2. Single-class networks

Consider an open system with a single class of individual (customer, call, or job). Individuals arrive as a Poisson process with state-dependent rate $\alpha(n)$, where n is the number of individuals currently in the system. Arriving individuals have *workloads* which are independent of each other and of the arrivals process with a common distribution μ on \mathbb{R}_+ which we assume to have a finite mean $m(\mu)$. While there are n individuals in the system, their *total* workload is reduced at a rate $\beta(n) \geq 0$, where we assume that $\beta(n) > 0$ if and only if $n > 0$; an individual

departs the system when its workload is reduced to zero. By suitably redefining the rates $\beta(n)$ if necessary, we may, and do, assume without loss of generality that the mean workload $m(\mu)$ is equal to 1.

We first consider the processor-sharing case. Here, when there are $n > 0$ individuals in the system, the workload of each is simultaneously reduced at a rate $\beta(n)/n$, and the setup described above becomes a fairly general description of a single-class processor-sharing system. A special case is the simple Erlang loss system, in which, for some $\alpha, \beta > 0$, we have $\alpha(n) = \alpha \mathbf{1}_{\{n < C\}}$ for some capacity $C \leq \infty$ (where $\mathbf{1}_{\{\cdot\}}$ is the indicator function) and $\beta(n) = n\beta$ for all $n \geq 0$. Here individuals are typically referred to as calls, and workloads correspond to call durations (since $\beta(n)/n$ is independent of n). A further special case is the M/GI/ m/∞ processor-sharing queue, in which, again for some $\alpha, \beta > 0$, we have $\alpha(n) = \alpha$ for all n and $\beta(n) = \min(n, m)\beta$ for all n and some fixed m .

We represent the system as a Markov process $(X(t))_{t \geq 0}$ by defining its state at any time t to be the number n of individuals then in the system together with their residual workloads at that time. (An alternative is to record, for each individual, the workload completed at time t .) For a given $n > 0$ these workloads form an (unordered) set, and may be regarded as taking values in the quotient space S_n obtained from \mathbb{R}_+^n by identifying points which may be obtained from each other under permutation of their coordinates. The σ -algebra $\mathcal{B}(S_n)$ on S_n is similarly formed in the obvious manner from the Borel σ -algebra on \mathbb{R}_+^n . The state space S for the process $(X(t))_{t \geq 0}$ is then the union of the $S_n, n \geq 0$, where the set S_0 is taken to consist of a single point, and its associated σ -algebra $\mathcal{B}(S)$ consists of those sets which are countable unions of sets in the σ -algebras $\mathcal{B}(S_n)$. The process $(X(t))_{t \geq 0}$ is thus an instance of a piecewise-deterministic Markov process (see Davis (1984), (1993)). However, we avoid the need for most of the general machinery for handling such processes.

We define the probability distribution $\bar{\mu}$ on \mathbb{R}_+ to be the stationary residual life distribution of the renewal process with interevent distribution μ , that is, if μ has distribution function F then $\bar{\mu}$ has distribution function G given by

$$G(x) = 1 - \int_x^\infty (1 - F(y)) \, dy$$

(recall that $m(\mu) = 1$). Note that the ‘residual life’ here should be thought of as a residual workload rather than a time. For each $n \geq 1$, define also the probability distribution $\bar{\mu}_n$ on S_n to be the product of n copies of the distribution $\bar{\mu}$, again with the above identification of points in \mathbb{R}_+^n (more formally, $\bar{\mu}_n(A) = \bar{\mu}^n(\theta^{-1}(A)), A \in \mathcal{B}(S_n)$, where $\bar{\mu}^n$ is the product of n copies of the distribution $\bar{\mu}$ and θ is the projection from \mathbb{R}_+^n into the quotient space S_n). Thus, $\bar{\mu}_n$ represents the joint distribution of the residual lives at any time in a set of n independent stationary renewal processes each with interevent distribution μ ; we also define $\bar{\mu}_0$ to be the probability distribution concentrated on the single-point set S_0 . For each n , we also regard $\bar{\mu}_n$ as a distribution on S , assigning its total mass of one to the set S_n . Finally, for any distribution π on \mathbb{Z}_+ , define the distribution $\bar{\mu}_\pi$ on S by

$$\bar{\mu}_\pi = \sum_{n \in \mathbb{Z}_+} \pi(n) \bar{\mu}_n.$$

Thus, $\bar{\mu}_\pi$ assigns probability $\pi(n)$ to the event that there are n individuals in the system and, conditional on this event, assigns the distribution $\bar{\mu}_n$ to their residual workloads.

Theorem 1. *Suppose that the distribution π on \mathbb{Z}_+ is the solution of the balance equations*

$$\pi(n + 1)\beta(n + 1) = \pi(n)\alpha(n), \quad n \geq 0, \tag{1}$$

and that

$$\sum_{n \geq 0} \pi(n)\alpha(n) < \infty. \tag{2}$$

Then the distribution $\bar{\mu}_\pi$ on S is stationary for the process $(X(t))_{t \geq 0}$ and, in particular, the distribution π is stationary for the associated number of individuals in the system.

Remark 1. The condition (2) ensures that, under stationarity, individuals arrive at the system at a finite rate.

Proof of Theorem 1. In order to exclude pathological behaviour in the argument below, we make the one additional assumption that the distribution μ has no atom of probability at zero. This is without loss of generality: in the case that μ does have such an atom, the evolution of the system may clearly be equivalently described by redefining α , β , and μ so as to remove it, and the result of the theorem is easily obtained via this reparametrisation.

Analogously to the definition of $\bar{\mu}_n$, for each $n \geq 1$, define the probability distribution $\hat{\mu}_n$ on S_n to be the product of $n - 1$ copies of the distribution $\bar{\mu}$ and a single copy of the distribution μ , yet again with the above identification of points in \mathbb{R}_+^n . (More formally, $\hat{\mu}_n(A) = \hat{\mu}^{(n)}(\theta^{-1}(A))$, $A \in \mathcal{B}(S_n)$, where $\hat{\mu}^{(n)} = \bar{\mu}^{n-1} \times \mu$ and θ is again the projection from \mathbb{R}_+^n into the quotient space S_n .) We again regard $\hat{\mu}_n$ as a distribution on S , assigning mass one to the set S_n .

Consider now the modified process $(\hat{X}(t))_{t \geq 0}$ on S describing the system in which the workload distribution is again μ and in which, when there are $n \geq 1$ individuals in the system, individual workloads are again reduced at rate $\beta(n)/n$. However, for the modified system, an individual departing on completion of its workload is immediately replaced by another bringing an independent workload with distribution μ , and external arrivals to the system are not accepted. Thus, for the modified system, the number of individuals remains constant, and, conditional on this being n , the system behaves as a set of n independent renewal processes, each of which has stationary residual workload distribution $\bar{\mu}$. Hence, for any distribution π' on \mathbb{Z}_+ , the distribution $\bar{\mu}_{\pi'}$ on S is stationary for the process $(\hat{X}(t))_{t \geq 0}$.

Let $(P_t)_{t \geq 0}$ and $(\hat{P}_t)_{t \geq 0}$ be the semigroups of transition kernels associated respectively with the processes $(X(t))_{t \geq 0}$ and $(\hat{X}(t))_{t \geq 0}$. For any $a > 0$, let \mathcal{D}_a be the class of functions f on S taking values in $[0, 1]$ and satisfying the continuity condition

$$|P_t f(x) - f(x)| \leq at, \quad \text{for all } x \in S \text{ and } t > 0, \tag{3}$$

where $P_t f(x) = \int_S P_t(x, dy) f(y)$. For any such f and for any distribution ν on S , define also $\nu f = \int_S f(x)\nu(dx)$ and, for any $t > 0$, define $\nu P_t f = \nu(P_t f)$ (so that $\nu P_t f$ is the expectation of $f(X(t))$ when $(X(t))_{t \geq 0}$ is given initial distribution ν); similarly define $\nu \hat{P}_t f$.

Now compare the behaviour of the processes $(X(t))_{t \geq 0}$ and $(\hat{X}(t))_{t \geq 0}$, each started with the distribution $\bar{\mu}_\pi$; so as to simplify the description below we couple these two processes so that they agree until the time of the first arrival or workload completion. We then find (see the further explanation below) that, with this common initial distribution, for any $a > 0$, $f \in \mathcal{D}_a$,

and $h > 0$,

$$\begin{aligned}
 \bar{\mu}_\pi P_h f - \bar{\mu}_\pi \hat{P}_h f &= \mathbb{E}(f(X(h)) - f(\hat{X}(h))) \\
 &= h \sum_{n \geq 0} \pi(n) [\alpha(n)(\hat{\mu}_{n+1} f - \bar{\mu}_n f) + \beta(n)(\bar{\mu}_{n-1} f - \hat{\mu}_n f)] + o(h) \quad (4) \\
 &= h \sum_{n \geq 0} [\pi(n)\alpha(n) - \pi(n+1)\beta(n+1)] (\hat{\mu}_{n+1} f - \bar{\mu}_n f) + o(h) \\
 &= o(h) \quad (5)
 \end{aligned}$$

as $h \rightarrow 0$ (recall in (4) that $\beta(0) = 0$). Furthermore, the above convergence as $h \rightarrow 0$ is uniform over $f \in \mathcal{D}_a$ in the sense that (5) may be written as

$$\sup_{f \in \mathcal{D}_a} |\bar{\mu}_\pi P_h f - \bar{\mu}_\pi \hat{P}_h f| = o(h) \quad \text{as } h \rightarrow 0 \quad (6)$$

(again see below). To show that (4) holds, note first that, from the above coupling and for any $h > 0$, we have $f(X(h)) = f(\hat{X}(h))$ except where there is either at least one external arrival or at least one workload completion in $[0, h]$. It follows from the definition of $\bar{\mu}_\pi$ that, conditional on the number of individuals initially being n , the probability of an external arrival in $[0, h]$ is $\alpha(n)h + o(h)$ as $h \rightarrow 0$, and that an arriving individual finding the distribution of the system to be $\bar{\mu}_n$ changes this to $\hat{\mu}_{n+1}$ in the case of the process $(X(t))_{t \geq 0}$ and leaves it unchanged in the case of the process $(\hat{X}(t))_{t \geq 0}$. Similarly, again conditional on the number of individuals initially being n (and recalling that $m(\mu) = 1$), the probability of a workload completion in a time interval $[0, h]$ is $\beta(n)h + o(h)$ as $h \rightarrow 0$, and, under the distribution $\bar{\mu}_n$, conditional on such a completion taking place, the residual workload distribution becomes $\bar{\mu}_{n-1}$ in the case of the process $(X(t))_{t \geq 0}$ and $\hat{\mu}_n$ in the case of the process $(\hat{X}(t))_{t \geq 0}$. Furthermore, it follows from (1) and (2) that, under the initial distribution $\bar{\mu}_\pi$, the probability of two or more arrivals or workload completions in $[0, h]$ is $o(h)$ as $h \rightarrow 0$. That the relation (4) now holds as $h \rightarrow 0$ with the uniformity over $f \in \mathcal{D}_a$ required for (6) follows easily from these results and from the definition of \mathcal{D}_a . To see this note that, since $f \in \mathcal{D}_a$ implies that f takes values in $[0, 1]$, the contribution to the error term in (4) resulting from the neglect of the possibility of two or more arrivals or workload completions in $[0, h]$ is uniformly $o(h)$ as $h \rightarrow 0$, as required. Similarly, the terms $\hat{\mu}_{n+1} f - \bar{\mu}_n f$ and $\bar{\mu}_{n-1} f - \hat{\mu}_n f$ in (4) are obtained by treating the precise time of the first arrival or workload completion within $[0, h]$ as if it were time h ; recalling that terms of type $\hat{\mu}_{n+1}$ are probability measures, it follows from (3) that the consequent error in each of the above two terms is bounded by $2ah$, so that the further contribution to the error term in (4) is $O(h^2)$ as $h \rightarrow 0$, again with uniformity over $f \in \mathcal{D}_a$. The relations (5) and, hence, (6) are now immediate from the balance equations (1). Since the distribution $\bar{\mu}_\pi$ is stationary for the process $(\hat{X}(t))_{t \geq 0}$, it now follows from (6) that, again for any $a > 0$ and $h > 0$,

$$\sup_{f \in \mathcal{D}_a} |\bar{\mu}_\pi P_h f - \bar{\mu}_\pi f| = o(h) \quad \text{as } h \rightarrow 0.$$

Furthermore, it is straightforward that if $f \in \mathcal{D}_a$ then $P_t f \in \mathcal{D}_a$ for any $t > 0$. Standard manipulations using the semigroup structure of $(P_t)_{t \geq 0}$, for example the consideration of increasingly refined partitions of the interval $[0, t]$, now give, for all $a > 0$, $f \in \mathcal{D}_a$, and $t \geq 0$,

$$\bar{\mu}_\pi P_t f = \bar{\mu}_\pi f. \quad (7)$$

Finally, we show that it follows from (7) that $\bar{\mu}_\pi P_t = \bar{\mu}_\pi$, for all $t > 0$, so that $\bar{\mu}_\pi$ is stationary for $(X(t))_{t \geq 0}$, as required. It is sufficient to show that, for any $n \geq 1$ and any set $A \in \mathcal{B}(S_n)$ whose inverse image in \mathbb{R}_+^n under the mapping θ defined above is a product of intervals in \mathbb{R}_+ , we have

$$\bar{\mu}_\pi P_t \mathbf{1}_A = \bar{\mu}_\pi \mathbf{1}_A. \tag{8}$$

It follows from the piecewise-deterministic form of the process $(X(t))_{t \geq 0}$ that we may choose a sequence of functions $(f_k, k \geq 1)$ such that, for each k ,

- (i) $f_k \in \mathcal{D}_a$ for some $a > 0$,
- (ii) f_k and $\mathbf{1}_A$ agree except on a set whose Lebesgue measure (under θ^{-1}) in \mathbb{R}_+^n tends to zero as $k \rightarrow \infty$.

Since $\bar{\mu}_\pi$ and, so, $\bar{\mu}_\pi P_t$ are nonatomic distributions, (8) now follows by using (7) with $f = f_k$ and letting $k \rightarrow \infty$.

Remark 2. Suppose that (1) is multiplied by the signed measure $\hat{\mu}_{n+1} - \bar{\mu}_n$ to give

$$\pi(n)\alpha(n)(\hat{\mu}_{n+1} - \bar{\mu}_n) = \pi(n+1)\beta(n+1)(\hat{\mu}_{n+1} - \bar{\mu}_n), \quad n \geq 0.$$

These equations have an obvious interpretation as representing, under the distribution $\bar{\mu}_\pi$ and for each $n \geq 0$, a detailed balance of flux between S_n and S_{n+1} , not just with regard to the total probability assigned to each of these spaces, but also with regard to the distribution of the residual workload sizes. This is the intuition underlying the derivation of (5) (which is also that of Pechinkin’s coupling approach), i.e. under $\bar{\mu}_\pi$, an arrival finding n individuals in the system transforms the residual workload distribution from $\bar{\mu}_n$ to $\hat{\mu}_{n+1}$, while a departure from the system when it contains $n + 1$ individuals transforms the residual workload distribution from what would have been $\hat{\mu}_{n+1}$, if the individual had remained in the system with a renewed workload, to the distribution $\bar{\mu}_n$.

In the case where we do not have processor sharing, i.e. in which it is no longer the case that at any time all workloads are being reduced at the same rate, it is necessary at any time to distinguish the individuals in the system. Thus, each S_n is replaced by \mathbb{R}_+^n and the state space S is replaced by $S^* = \bigcup_{n \geq 0} \mathbb{R}_+^n$. We consider as an example the case of the single-server queue with ‘last-in–first-out preemptive-resume’ discipline, in which at any time all service effort is devoted to the last individual to arrive at the system. If at any time there are n individuals in the system, we may index these by $i = 1, \dots, n$ in the order of their arrival, and no individual changes index during its time in the system; as usual arrivals occur as a Poisson process with rate $\alpha(n)$, and the workload of individual n is now being reduced at rate $\beta(n)$, while that of the remaining individuals is being reduced at rate zero. As previously, define the probability distribution $\bar{\mu}$ on \mathbb{R}_+ to be the stationary residual life distribution of the renewal process with interevent distribution μ and, for each $n \geq 0$, let the distribution $\bar{\mu}_n$ on \mathbb{R}_+^n now be the (ordered) product of n copies of $\bar{\mu}$. For each $n \geq 1$, let the distribution $\hat{\mu}_n$ on \mathbb{R}_+^n be the (ordered) product of $n - 1$ copies of the distribution $\bar{\mu}$ and a single copy of the distribution μ , with the latter assigned to the n th coordinate of \mathbb{R}_+^n . Finally, for any distribution π on \mathbb{Z}_+ , define the distribution $\bar{\mu}_\pi$ on S by $\bar{\mu}_\pi = \sum_{n \in \mathbb{Z}_+} \pi(n) \bar{\mu}_n$, as before. With these (re)definitions, both Theorem 1 and its proof remain unchanged as stated. Again the underlying reason is as given in Remark 2, i.e. under the distribution $\bar{\mu}_\pi$, and relative to the modified process considered in the proof of Theorem 1, an arrival finding n individuals in the system transforms the residual workload distribution from

$\bar{\mu}_n$ to $\hat{\mu}_{n+1}$, while a departure from the system when it contains $n + 1$ individuals transforms the residual workload distribution from $\hat{\mu}_{n+1}$ to $\bar{\mu}_n$. Note that this balance does not obtain in the case of, for example, a ‘first-in–first-out’ discipline, and here, as is again well known, we do not have the above insensitivity.

3. Multi-class networks

Now consider a multi-class network. We concentrate on the processor-sharing case – adaptations to other disciplines may be made as in the single-class case. Let $\mathcal{I} = \{1, \dots, N\}$ denote the set of classes and let $\mathbf{n} = (n_i, i \in \mathcal{I})$, where n_i is the number of individuals in each class i . An individual entering class i acquires a workload which has distribution μ^i with nonzero finite mean $m(\mu^i)$; we again assume without loss of generality that $m(\mu^i) = 1$. The workload of each individual in class i is reduced at a state-dependent rate $\phi_i(\mathbf{n})/n_i$, where $\phi_i(\mathbf{n}) > 0$ if and only if $n_i > 0$. Individuals arrive at each class i from outside the network as a Poisson process with state-dependent rate $\phi_{0i}(\mathbf{n})$; on completion of its workload in any class i an individual moves to class j with state-dependent probability $\phi_{ij}(\mathbf{n})/\phi_i(\mathbf{n})$ or leaves the network with probability $\phi_{i0}(\mathbf{n})/\phi_i(\mathbf{n})$, where

$$\sum_{j \in \mathcal{I}} \phi_{ij}(\mathbf{n}) + \phi_{i0}(\mathbf{n}) = \phi_i(\mathbf{n}) \tag{9}$$

(there are no problems in allowing the possibility $\phi_{ii}(\mathbf{n}) > 0$). The workloads, arrival processes, and routing decisions are all independent.

As in the single-class case, we represent the system as a Markov process $(X(t))_{t \geq 0}$ by defining its state at any time to be the vector \mathbf{n} introduced above together with the residual workloads at that time of the set of individuals in each class. For given \mathbf{n} , these workloads take values in the space $S_{\mathbf{n}}$, which is the ordered product of the spaces S_{n_i} , $i \in \mathcal{I}$, where, as previously, each S_{n_i} is formed from $\mathbb{R}_+^{n_i}$ by identifying points within the latter space which may be obtained from each other under permutation of their coordinates (and where again the set S_0 contains a single point). The state space S for the system is the union of all the possible $S_{\mathbf{n}}$, and the spaces $S_{\mathbf{n}}$ and S are endowed with the obvious σ -algebras $\mathcal{B}(S_{\mathbf{n}})$ and $\mathcal{B}(S)$.

Analogously to the single-class case, for each $i \in \mathcal{I}$ define the probability distribution $\bar{\mu}^i$ on \mathbb{R}_+ to be the stationary residual life distribution of the renewal process with interevent distribution μ^i (as previously, the residual life should be interpreted as a residual workload). For each $i \in \mathcal{I}$ and for each $n_i \geq 1$, define as previously the distribution $\bar{\mu}_{n_i}^i$ on S_{n_i} to be the product of n_i copies of the distribution $\bar{\mu}^i$ (again with the above identification of points in $\mathbb{R}_+^{n_i}$) – representing the joint distribution of the residual lives in a set of n_i independent stationary renewal processes each with interevent distribution μ^i . Also, define $\bar{\mu}_0^i$ to be the probability distribution concentrated on the single-point set S_0 . For each $\mathbf{n} \in \mathbb{Z}_+^N$, define the distribution $\bar{\mu}_{\mathbf{n}}$ on $S_{\mathbf{n}}$ to be the (ordered) product distribution which, for each $i \in \mathcal{I}$, assigns the distribution $\bar{\mu}_{n_i}^i$ to S_{n_i} . We again regard the distribution $\bar{\mu}_{\mathbf{n}}$ as a distribution on S , assigning its total mass of one to the set $S_{\mathbf{n}}$. For any positive distribution π on \mathbb{Z}_+^N , we define the distribution $\bar{\mu}_{\pi}$ on S by $\bar{\mu}_{\pi} = \sum_{\mathbf{n} \in \mathbb{Z}_+^N} \pi(\mathbf{n}) \bar{\mu}_{\mathbf{n}}$.

It is notationally convenient to expand the set \mathcal{I} to $\mathcal{I}' = \{0\} \cup \mathcal{I}$, treating 0 as an extra class feeding external arrivals to, and receiving departures from, the network. (However, the components of the state \mathbf{n} of the network remain indexed in the original set \mathcal{I} .) For completeness, we define $\phi_{00}(\mathbf{n}) = 0$, for all \mathbf{n} , and $\phi_{ij}(\mathbf{n}) = 0$, for all $i \in \mathcal{I}$, $j \in \mathcal{I}'$, and \mathbf{n} such that $n_i = 0$

(so that (9) also remains valid for such \mathbf{n}). For each $i \in \mathcal{I}$, let \mathbf{e}_i be the N -dimensional vector whose i th component is 1 and whose other components are 0, and let \mathbf{e}_0 be the N -dimensional vector all of whose components are 0. For each \mathbf{n} and each $i, j \in \mathcal{I}'$, define the vectors $T_i^j \mathbf{n} = \mathbf{n} - \mathbf{e}_i + \mathbf{e}_j$, $T_i \mathbf{n} = T_i^0 \mathbf{n} = \mathbf{n} - \mathbf{e}_i$, and $T^j \mathbf{n} = T_0^j \mathbf{n} = \mathbf{n} + \mathbf{e}_j$.

Theorem 2. *Suppose that the distribution π on \mathbb{Z}_+^N satisfies the partial balance equations*

$$\pi(\mathbf{n}) \sum_{j \in \mathcal{I}'} \phi_{ij}(\mathbf{n}) = \sum_{j \in \mathcal{I}'} \pi(T_i^j \mathbf{n}) \phi_{ji}(T_i^j \mathbf{n}), \quad \mathbf{n} \in \mathbb{Z}_+^N, \quad i \in \mathcal{I}', \tag{10}$$

where, for \mathbf{n} and $i \in \mathcal{I}$ such that $n_i = 0$, we interpret the right-hand side of (10) as zero (recall that when $n_i = 0$ we have $\phi_{ij}(\mathbf{n}) = 0$ for all $j \in \mathcal{I}'$, so that (10) is automatically satisfied in this case). Suppose also that

$$\sum_{\mathbf{n} \in \mathbb{Z}_+^N} \pi(\mathbf{n}) \sum_{i \in \mathcal{I}} \phi_{0i}(\mathbf{n}) < \infty. \tag{11}$$

Then $\bar{\mu}_\pi$ is stationary for the process $(X(t))_{t \geq 0}$ and, in particular, π is stationary for the associated numbers of individuals in the system. Conversely, if a distribution π on \mathbb{Z}_+^N is stationary for the numbers of individuals in the system for all $\boldsymbol{\mu} = (\mu^i, i \in \mathcal{I})$ such that $m(\mu^i) = 1$ for all $i \in \mathcal{I}$, then π satisfies (10).

Proof. Suppose first that π satisfies (10). As in the proof of Theorem 1, we again assume without loss of generality that each distribution μ^i has no atom of probability at zero.

For each \mathbf{n} and for each i such that $n_i \geq 1$, also define the residual workload distribution $\hat{\mu}_\mathbf{n}^i$ on $S_\mathbf{n}$ by

$$\hat{\mu}_\mathbf{n}^i = \hat{\mu}_{n_i}^i \prod_{j \neq i} \bar{\mu}_{n_j}^j,$$

where $\hat{\mu}_{n_i}^i$ is defined as in the proof of Theorem 1. Thus, $\hat{\mu}_\mathbf{n}^i$ corresponds to each individual in each class j independently having the stationary residual workload distribution $\bar{\mu}^j$, except only that a single individual in the class i is given the workload distribution μ^i . For each \mathbf{n} , also define $\hat{\mu}_\mathbf{n}^0 = \bar{\mu}_\mathbf{n}$.

Again, as in the proof of Theorem 1, define the process $(\hat{X}(t))_{t \geq 0}$ on S to be that appropriate to the modified system in which there are no arrivals, departures, or transfers between classes; rather each individual in each class i , on completion of its workload, acquires a new independent workload with distribution μ^i . Thus, the occupancy of the system remains constant; conditional on this being \mathbf{n} , individual workloads in any class i such that $n_i > 1$ are again reduced at rate $\phi_i(\mathbf{n})/n_i$ and the system behaves as a set of independent renewal processes. Furthermore, for any distribution π' on \mathbb{Z}_+^N , the distribution $\bar{\mu}_{\pi'}$ on S is stationary for $(\hat{X}(t))_{t \geq 0}$.

Again, let $(P_t)_{t \geq 0}$ and $(\hat{P}_t)_{t \geq 0}$ be the semigroups of transition kernels associated, respectively, with the processes $(X(t))_{t \geq 0}$ and $(\hat{X}(t))_{t \geq 0}$ and, for any $a > 0$, let \mathcal{D}_a be the class of functions f on S taking values in $[0, 1]$ and satisfying the earlier continuity condition (3). Comparison of the behaviours of the processes $(X(t))_{t \geq 0}$ and $(\hat{X}(t))_{t \geq 0}$, each started with the distribution $\bar{\mu}_\pi$ and coupled as in the earlier proof until the time of the first external arrival or

workload completion, now gives, for any $a > 0$, $f \in \mathcal{D}_a$, and $h > 0$,

$$\begin{aligned}
 & \bar{\mu}_\pi P_h f - \bar{\mu}_\pi \hat{P}_h f \\
 &= E(f(X(h)) - f(\hat{X}(h))) \\
 &= h \sum_{\mathbf{n} \in \mathbb{Z}_+^N} \pi(\mathbf{n}) \sum_{i \in \mathcal{I}'} \sum_{j \in \mathcal{I}'} \phi_{ij}(\mathbf{n}) (\hat{\mu}_{T_i^j \mathbf{n}}^j f - \hat{\mu}_\mathbf{n}^i f) + o(h) \tag{12} \\
 &= h \sum_{\mathbf{n} \in \mathbb{Z}_+^N} \left(\sum_{i \in \mathcal{I}'} \sum_{j \in \mathcal{I}'} \pi(\mathbf{n}) \phi_{ji}(\mathbf{n}) \hat{\mu}_{T_j^i \mathbf{n}}^i f - \sum_{i \in \mathcal{I}'} \sum_{j \in \mathcal{I}'} \pi(\mathbf{n}) \phi_{ij}(\mathbf{n}) \hat{\mu}_\mathbf{n}^i f \right) + o(h) \\
 &= h \sum_{\mathbf{n} \in \mathbb{Z}_+^N} \left(\sum_{i \in \mathcal{I}'} \sum_{j \in \mathcal{I}'} \pi(T_i^j \mathbf{n}) \phi_{ji}(T_i^j \mathbf{n}) \hat{\mu}_\mathbf{n}^i f - \sum_{i \in \mathcal{I}'} \sum_{j \in \mathcal{I}'} \pi(\mathbf{n}) \phi_{ij}(\mathbf{n}) \hat{\mu}_\mathbf{n}^i f \right) + o(h) \\
 &= h \sum_{\mathbf{n} \in \mathbb{Z}_+^N} \sum_{i \in \mathcal{I}'} \left(\sum_{j \in \mathcal{I}'} \pi(T_i^j \mathbf{n}) \phi_{ji}(T_i^j \mathbf{n}) - \sum_{j \in \mathcal{I}'} \pi(\mathbf{n}) \phi_{ij}(\mathbf{n}) \right) \hat{\mu}_\mathbf{n}^i f + o(h) \\
 &= o(h) \tag{13}
 \end{aligned}$$

as $h \rightarrow 0$, with uniformity of convergence over all $f \in \mathcal{D}_a$. We may therefore write

$$\sup_{f \in \mathcal{D}_a} |\bar{\mu}_\pi P_h f - \bar{\mu}_\pi \hat{P}_h f| = o(h) \quad \text{as } h \rightarrow 0 \tag{14}$$

(recall again that $\phi_{ij}(\mathbf{n}) = 0$ whenever $n_i = 0$, so that there is no difficulty with the lack of a formal definition of $\hat{\mu}_\mathbf{n}^i$ in this case). The identity (14) is simply the multi-class version of the identity (6) in the proof of Theorem 1, and is similarly obtained, albeit with a slightly more compact notation. In particular, conditional on the common initial distribution of the two processes being given by $\bar{\mu}_\pi$, in the time interval $[0, h]$ (where h is small), a transition from i to j in the original system – where either i or j may be zero – occurs with probability $\phi_{ij}(\mathbf{n})h + o(h)$ as $h \rightarrow 0$, and in this case the distribution of the process $(X(t))_{t \geq 0}$ becomes $\hat{\mu}_{T_i^j \mathbf{n}}^j$ while that of the process $(\hat{X}(t))_{t \geq 0}$ becomes $\hat{\mu}_\mathbf{n}^i$. It follows that (12) now holds (with the required uniformity of convergence), also using (11), as in the earlier proof; finally the results (13) and, so, (14) follow from the partial balance equations (10). Since the distribution $\bar{\mu}_\pi$ is stationary for the process $(\hat{X}(t))_{t \geq 0}$, it now follows from (14) that, again for any $a > 0$ and $h > 0$,

$$\sup_{f \in \mathcal{D}_a} |\bar{\mu}_\pi P_h f - \bar{\mu}_\pi f| = o(h) \quad \text{as } h \rightarrow 0.$$

It now follows that $\bar{\mu}_\pi$ is stationary for $(X(t))_{t \geq 0}$, as in the proof of Theorem 1.

Now suppose that a distribution π on \mathbb{Z}_+^N is stationary for the numbers of individuals in the system for all $\boldsymbol{\mu} = (\mu^i, i \in \mathcal{I})$ with $m(\mu^i) = 1$ for all $i \in \mathcal{I}$. A proof that π then necessarily satisfies the partial balance equations (10) was given by Bonald and Proutière (2002), (2003). In summary, consider the case in which in every class the workload distribution is exponential with mean one and, for any fixed class i , compare this with the case in which, for some λ , $0 < \lambda < 1$, the workload distribution in class i is replaced by a mixture of two distributions, obtained by choosing with probability λ an exponential distribution with mean λ^{-1} and with probability $1 - \lambda$ the distribution concentrated on zero. Both these models may be (re)formulated as simple Markov jump processes; in the latter case the transition rates into and out of the class i are reduced by a factor λ . Since π is stationary in both cases, comparison of the (full) balance equations for stationarity yields the partial balance equations (10).

Example 1. (*Processor-sharing Whittle networks.*) Suppose that, for some $\nu \geq 0$, some strictly positive function Φ on \mathbb{Z}_+^N , and some stochastic matrix $\mathbf{P} = (p_{ij}, i \in \mathcal{I}', j \in \mathcal{I}')$ such that $p_{00} = 0$, we have, for each \mathbf{n} ,

$$\begin{aligned} \phi_{0j}(\mathbf{n}) &= \nu p_{0j}, & j \in \mathcal{I}', \\ \phi_{ij}(\mathbf{n}) &= \frac{\Phi(T_i \mathbf{n})}{\Phi(\mathbf{n})} p_{ij}, & i \in \mathcal{I}, j \in \mathcal{I}', \end{aligned}$$

where we again use the convention that $\Phi(T_i \mathbf{n}) = 0$ whenever $n_i = 0$. Then it is readily checked that the partial balance equations (10) are satisfied by

$$\pi(\mathbf{n}) = a \Phi(\mathbf{n}) \prod_{i \in \mathcal{I}} \rho_i^{n_i}, \tag{15}$$

for any $a > 0$ and positive solution $\boldsymbol{\rho} = (\rho_i, i \in \mathcal{I})$ of the equations

$$\begin{aligned} \nu &= \sum_{j \in \mathcal{I}} \rho_j p_{j0}, \\ \rho_i &= \sum_{j \in \mathcal{I}} \rho_j p_{ji} + \nu p_{0i}, & i \in \mathcal{I}. \end{aligned}$$

Thus, in particular, the stationary distribution π given by (15) for the number of individuals of each type in the system is insensitive to the μ^i (recall our assumption that $m(\mu^i) = 1$ for all i). For the case where \mathbf{P} is irreducible and $\nu > 0$, the above equations for $\boldsymbol{\rho}$ have a unique solution. Now, when \mathbf{P} is irreducible and when $\nu = 0$ (corresponding to a closed network), π remains uniquely determined, up to a multiplicative constant, by (15). The case where $\Phi(\mathbf{n}) = \prod_{i \in \mathcal{I}} \lambda_i^{n_i}$ for positive constants $(\lambda_i, i \in \mathcal{I})$ characterises processor-sharing Jackson networks. Further discussion of Whittle networks was given by Serfozo (1999) and, for processor-sharing networks, by Bonald and Proutière (2002), (2003).

Example 2. (*Networks with no internal transitions.*) Suppose that $\phi_{ij}(\mathbf{n}) = 0$ for all $i, j \in \mathcal{I}$ and for all $\mathbf{n} \in \mathbb{Z}_+^N$, so that no transitions are possible between the classes in \mathcal{I} . The partial balance equations (10) then reduce to the detailed balance equations

$$\pi(\mathbf{n}) \phi_{i0}(\mathbf{n}) = \pi(T_i \mathbf{n}) \phi_{0i}(T_i \mathbf{n}), \quad \mathbf{n} \in \mathbb{Z}_+^N, n_i \geq 1, i \in \mathcal{I}. \tag{16}$$

(In the case of a single class, these equations further reduce to (1).) An example is given by a traditional (uncontrolled) loss network; see, for example, Kelly (1986). This is naturally processor-sharing. Here, workloads are identified with call durations and, for some set $\mathcal{A} \subset \mathbb{Z}_+^N$ such that $\mathbf{n} \in \mathcal{A}$ implies that $T_i \mathbf{n} \in \mathcal{A}$ for all \mathbf{n} and i such that $n_i \geq 1$ (\mathcal{A} is typically defined by capacity constraints), we have

$$\begin{aligned} \phi_{0i}(\mathbf{n}) &= \nu_i \mathbf{1}_{\{T^i \mathbf{n} \in \mathcal{A}\}}, \\ \phi_{i0}(\mathbf{n}) &= \sigma_i n_i, \end{aligned}$$

for some vectors $(\nu_i, i \in \mathcal{I})$ and $(\sigma_i, i \in \mathcal{I})$ of strictly positive parameters. The equations (16) are then satisfied by

$$\pi(\mathbf{n}) = a \prod_{i \in \mathcal{I}} \frac{\kappa_i^{n_i}}{n_i!},$$

where $\kappa_i = v_i/\sigma_i$ for each i , and where a is naturally chosen to be a normalising constant. As was originally shown by Burman *et al.* (1984), we therefore again have insensitivity of the occupancy distribution π of the network. The stationary distribution of the residual call durations is as identified by Theorem 2.

Other examples of processor-sharing networks with no internal transitions are given by those used to model connections in communications networks with simultaneous resource requirements and variable bandwidth requirements; see, for example, Bonald and Massoulié (2001) and De Veciana *et al.* (2001). Here it is far from automatic that the detailed balance equations (16) are satisfied.

Acknowledgements

The author is most grateful to Serguei Foss, Takis Konstantopoulos, and Ilze Ziedins for some very helpful discussions, and to the referee for a careful reading of the manuscript and some most helpful suggestions.

References

- BONALD, T. AND MASSOULIÉ, L. (2001). Impact of fairness on Internet performance. In *Proc. ACM SIGMETRICS 2001* (Cambridge, MA), ACM Press, New York.
- BONALD, T. AND PROUTIERE, A. (2002). Insensitivity in processor-sharing networks. *Performance Evaluation* **49**, 193–209.
- BONALD, T. AND PROUTIERE, A. (2003). Insensitive bandwidth sharing in data networks. *Queueing Systems* **44**, 69–100.
- BURMAN, D. Y. (1981). Insensitivity in queueing systems. *Adv. Appl. Prob.* **13**, 846–859.
- BURMAN, D. Y., LEHOCZKY, J. P. AND LIM, Y. (1984). Insensitivity of blocking probabilities in a circuit switching network. *J. Appl. Prob.* **21**, 850–859.
- DAVIS, M. H. A. (1984). Piecewise-deterministic Markov processes: a general class of non-diffusion stochastic models. *J. R. Statist. Soc. B* **46**, 353–388.
- DAVIS, M. H. A. (1993). *Markov Models and Optimization*. Chapman & Hall, London.
- DE VECIANA, G., LEE, T.-J. AND KONSTANTOPOULOS, T. (2001). Stability and performance analysis of networks supporting elastic services. *IEEE/ACM Trans. Networking* **9**, 2–14.
- KELLY, F. P. (1979). *Reversibility and Stochastic Networks*. John Wiley, Chichester.
- KELLY, F. P. (1986). Blocking probabilities in large circuit-switched networks. *Adv. Appl. Prob.* **18**, 473–505.
- PECHINKIN, A. V. (1983). An invariant queueing system. *Math. Operationsforsch. Statist. Ser. Optim.* **14**, 433–444.
- PECHINKIN, A. V. (1987). A new proof of Erlang's formula for a lossy multichannel queueing system. *Soviet J. Comput. System Sci.* **25**, 165–168.
- SCHASSBERGER, R. (1978). Insensitivity of steady-state distributions of generalized semi-Markov processes with speeds. *Adv. Appl. Prob.* **10**, 836–851.
- SCHASSBERGER, R. (1986). Two remarks on insensitive stochastic models. *Adv. Appl. Prob.* **18**, 791–814.
- SERFOZO, R. F. (1999). *Introduction to Stochastic Networks*. Springer, Berlin.
- SEVAST'YANOV, B. A. (1957). An ergodic theory for Markov processes and its application to a telephone system with refusals. *Theory Prob. Appl.* **2**, 104–112.
- WHITTLE, P. (1985). Partial balance and insensitivity. *J. Appl. Prob.* **22**, 168–176.