

STABILITY OF PROCESSOR SHARING NETWORKS WITH SIMULTANEOUS RESOURCE REQUIREMENTS

JENNIE HANSEN,* **

CIAN REYNOLDS* AND

STAN ZACHARY,* *Heriot-Watt University*

Abstract

We study the phenomenon of *entrainment* in processor sharing networks, whereby, while individual network resources have sufficient capacity to meet demand, the requirement for simultaneous availability of resources means that a network may nevertheless be unstable. We show that instability occurs through poor control, and that, for a variety of network topologies, only small modifications to controls are required in order to ensure stability. For controls which possess a natural monotonicity property, we give some new results for the classification of the corresponding Markov processes, which lead to conditions both for stability and for instability.

Keywords: Processor sharing network; stability; entrainment

2000 Mathematics Subject Classification: Primary 60K20

Secondary 60K25

1. Introduction

Modern communications networks, such as the Internet, are able at any time to share their resources, for example bandwidth, among those calls or connections currently in progress. Such calls may require simultaneous capacity from several resources in the network. For example, ‘streaming’ applications require, for their duration, a reservation of bandwidth on each network link over which they connect. It may then happen that while each resource in the network, considered in isolation, has sufficient capacity to service the demand placed on it, the control of the network is such that the requirement for *simultaneous* availability of capacity ensures that over time demand cannot be met, and that the network is unstable, that is, that the number of calls present in the network tends to ∞ . This is the phenomenon of *entrainment* and has been previously studied in this context by various authors, notably Bonald and Massoulié (2001), de Veciana *et al.* (2001), and Kelly and Williams (2004). In particular, these authors have considered a broad class of so-called *fair-sharing* control strategies; see below. Bonald and Massoulié and de Veciana *et al.* have shown that under any fair-sharing control the instability problems referred to above do not arise, i.e. provided that the various network resources have sufficient capacity individually, the network will remain stable. In the present paper we study the phenomenon of entrainment in more detail and characterise some of those conditions which give rise to its occurrence. In particular, we show that entrainment frequently arises through poor control whenever there are no calls of certain classes in the

Received 23 May 2006; revision received 15 May 2007.

* Postal address: Actuarial Mathematics and Statistics, School of Mathematical and Computer Sciences, Heriot-Watt University, Edinburgh EH14 4AS, UK.

** Email address: jhansen@ma.hw.ac.uk

network, and that only minor adjustments to control strategies are required in order to avoid it. Thus, very flexible management schemes, including those utilising significant prioritisation, may be safely implemented. We further establish, in Section 3, some new results for the stability of Markov chains whose transition rates possess a natural monotonicity property, yielding both necessary and sufficient conditions for the stability of a wide class of network controls. Since network parameters, in particular call arrival rates, are unlikely to be known in advance, it is of particular concern to identify controls whose stability is robust with respect to variations in these parameters.

We take as our model the following, which is essentially that introduced by Roberts and Massoulié (2000). Let R denote the finite set of possible call, or connection, types. We denote the state of the network at time t by $\mathbf{n}(t) = (n_r(t), r \in R)$, where $n_r(t)$ is the number of calls of each type r in progress at that time. Calls of each type $r \in R$ arrive at the network as a Poisson process with rate ν_r and have sizes which are exponentially distributed with mean μ_r^{-1} . Arrival processes and call sizes are all independent. (As usual, the above distributional assumptions make for simplicity of analysis. However, it seems likely that all the results of this paper remain qualitatively correct for other distributions of interarrival times and call sizes, subject only to these distributions having finite means, and to the same independence assumptions.) When the state of the network is \mathbf{n} , calls of each type r are allocated in total a bandwidth $b_r(\mathbf{n}) \geq 0$. We always assume that $b_r(\mathbf{n}) = 0$ whenever $n_r = 0$. The process $\mathbf{n}(\cdot)$ is thus Markov with state space $\mathbb{Z}_+^{|R|}$ and transition rates given by, for all $r \in R$,

$$\mathbf{n} \rightarrow \begin{cases} \mathbf{n} + \mathbf{e}_r & \text{at rate } \nu_r, \\ \mathbf{n} - \mathbf{e}_r & \text{at rate } \mu_r b_r(\mathbf{n}), \end{cases}$$

where \mathbf{e}_r denotes the vector whose r th component is 1 and whose other components are 0. For each \mathbf{n} , define also $\mathbf{b}(\mathbf{n}) = (b_r(\mathbf{n}), r \in R)$. We shall refer to $\mathbf{b} = (\mathbf{b}(\mathbf{n}), \mathbf{n} \in \mathbb{Z}_+^{|R|})$ as the control for the Markov process $\mathbf{n}(\cdot)$. The set of feasible controls \mathbf{b} is defined by a set of capacity constraints

$$\sum_{r \in R} a_{jr} b_r(\mathbf{n}) \leq c_j, \quad j \in J, \tag{1}$$

indexed in a finite set J , where each $a_{jr} \in \{0, 1\}$. Here c_j may be thought of as the capacity of resource j , and a call of type r utilises resource j if and only if $a_{jr} = 1$.

In order to allow some results to be stated with sufficient generality, it is convenient to allow the possibility that, for any j , we may have $c_j = \infty$ (corresponding to the effective nonexistence of the resource constraint j). However, we assume, without loss of generality, that, for all $r \in R$, there exists some $j \in J$ with $c_j < \infty$ and $a_{jr} = 1$.

Our interest is in identifying, for fixed values of the parameters ν_r, μ_r, c_j , and $a_{jr}, r \in R, j \in J$ (which we shall regard as defining a given network) those controls \mathbf{b} which are both feasible and stable, where we now take the latter to mean that the corresponding Markov process $\mathbf{n}(\cdot)$ is positive recurrent. A control \mathbf{b} for which $\mathbf{n}(\cdot)$ is null recurrent or transient is referred to as unstable.

For any feasible control \mathbf{b} , for any \mathbf{n} , we shall say that a resource $j \in J$ is saturated if the corresponding constraint, (1), is satisfied with equality. Following Bonald and Massoulié (2001), we shall further say that a feasible control \mathbf{b} is Pareto efficient if, for all \mathbf{n} and for all r such that $n_r > 0$, there exists $j \in J$ with $a_{jr} = 1$ such that j is saturated (so that $b_r(\mathbf{n})$ may not be increased without either decreasing $b_{r'}(\mathbf{n})$ for some $r' \neq r$ or else violating the constraints, (1)). It is sometimes helpful to consider controls which are not feasible, so we note in particular that the requirement of Pareto efficiency includes that of feasibility.

For each $r \in R$ we define $\kappa_r = \nu_r/\mu_r$, which may be thought of as the rate at which ‘work’ of type r arrives at the network. Many, but not all, stability results depend on the parameters ν_r and μ_r only through the corresponding κ_r . As observed by Bonald and Massoulié (2001), a necessary and sufficient condition for the existence of a stable feasible control is given by

$$\sum_{r \in R} a_{jr} \kappa_r < c_j, \quad j \in J. \tag{2}$$

For the necessity of this condition, observe that if it is violated for some j , then since, for any feasible control \mathbf{b} , we then have $\sum_{r \in R} a_{jr} (\kappa_r - b_r(\mathbf{n})) \geq 0$, standard arguments—see, for example, Asmussen (2003, Chapter 1, Proposition 5.4) or the argument of Kelly and Williams (2004)—show that \mathbf{b} cannot be stable. For the sufficiency of condition (2), define, for any vector $\hat{\mathbf{b}} = (\hat{b}_r, r \in R)$, the corresponding *complete partitioning* control \mathbf{b} by $b_r(\mathbf{n}) = \hat{b}_r$ whenever $n_r > 0$. Note that this defines a process $\mathbf{n}(\cdot)$ which corresponds to $|R|$ independent queues and is such that, in each queue r , arrivals occur at rate ν_r and departures (when $n_r > 0$) occur at rate $\mu_r \hat{b}_r$. Thus, a necessary and sufficient condition for the stability of this control is that $\hat{b}_r > \kappa_r$ for all $r \in R$. In particular, we may choose $\hat{\mathbf{b}}$ such that the corresponding complete partitioning control is stable and feasible if and only if condition (2) holds. Furthermore, given such a $\hat{\mathbf{b}}$, we may clearly define a Pareto efficient control \mathbf{b}' such that $b'_r(\mathbf{n}) \geq \hat{b}_r$ for all \mathbf{n} and for all r such that $n_r > 0$. The corresponding process $\mathbf{n}'(\cdot)$ may then be coupled to the process $\hat{\mathbf{n}}(\cdot)$ corresponding to $\hat{\mathbf{b}}$ in such a way that $n'_r(t) \leq \hat{n}_r(t)$ for all r and for all $t \geq 0$. Hence, condition (2) is also sufficient for the existence of *some* stable Pareto efficient control.

It is also clear that, in the case $|J| = 1$ of a single resource constraint, condition (2) is sufficient to ensure the stability of any Pareto efficient control. This is not in general true when $|J| > 1$, as is shown by Example 1 below, which is a simplification of one given by Bonald and Massoulié (2001).

Example 1. Suppose that $R = \{1, 2\}$, $J = \{1, 2\}$, and that the matrix $\mathbf{A} = (a_{jr}, j \in J, r \in R)$ is given by

$$\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

Thus, calls of type 1 are constrained by resource 1 only, while calls of type 2 are constrained by both resources 1 and 2. Suppose, without loss of generality, that $c_2 \leq c_1$. We assume condition (2), which here becomes

$$\kappa_1 + \kappa_2 < c_1, \quad \kappa_2 < c_2. \tag{3}$$

Consider the Pareto efficient control in which complete priority is given to calls of type 1. Then the long-run fraction of time in which the network is empty of calls of type 1—and so resource 1 is available for use by calls of type 2—is given by $1 - \kappa_1/c_1$. Since, when this resource is available, calls of type 2 are processed at rate c_2 , standard arguments for the stability of a single-server queue now show that the control is stable if and only if $\kappa_2 < c_2(1 - \kappa_1/c_1)$. This is a condition which is generally more restrictive than (3), above. When it is violated we have the phenomenon of entrainment referred to above; that is, from (3), each resource in the network, considered in isolation, has sufficient capacity, but the given Pareto efficient control is nevertheless unstable.

In Example 1 instability may be considered as occurring because such high priority is given to calls of type 1 as to ensure that, when these are emptied from the system, resource 1 is

thereafter consistently underutilized. As we shall show later, this problem may in general be avoided by choosing Pareto efficient controls \mathbf{b} such that, for each r , $b_r(\mathbf{n})$ is usually small whenever n_r is small. In particular, this property is possessed by the class of (weighted) fair-sharing controls introduced by Mo and Walrand (2000) in a generalisation of various classes considered earlier; see, for example, Kelly *et al.* (1998). These are defined as follows: for $\alpha > 0$, $\alpha \neq 1$, and weights w_r , $r \in R$, a weighted α -fair control, or bandwidth allocation, is given by taking, for each \mathbf{n} , $\mathbf{b}(\mathbf{n})$ to maximise the concave function

$$\sum_{r \in R} w_r n_r^\alpha \frac{b_r(\mathbf{n})^{1-\alpha}}{1-\alpha}, \tag{4}$$

subject to the capacity constraints, (1). This class is further extended to each of the cases $\alpha = 0, 1, \infty$ by taking the limit of the α -fair control as α tends to each of these values. (For the case in which $\alpha = 1$, this is equivalent to replacing the quantity defined in (4) by $\sum_{r \in R} w_r n_r \log b_r(\mathbf{n})$.) For the cases in which $\alpha = 1$ and $\alpha = \infty$, De Veciana *et al.* (2001) used Lyapunov function techniques to show that condition (2) is sufficient for the stability of (weighted) α -fair controls. Bonald and Massoulié (2001) showed that this result holds for general α by using fluid limits and appealing to a result of Dai (1995) for multiclass queueing networks. (In fact it is not certain that Dai’s result is directly applicable to the present networks with their simultaneous resource requirements; however, Bonald and Massoulié’s approach is essentially equivalent to showing that their function f defined by

$$f(\mathbf{n}) = \sum_{r \in R} w_r \mu_r^{-1} \kappa_r^{-\alpha} \frac{n_r^{\alpha+1}}{\alpha+1}$$

is a suitable Lyapunov function for establishing the sufficiency of (2) for the stability of α -fair controls, and this requires only a small extension to their argument.) We note, in particular, that the stability of the fair-sharing controls is robust with respect to parameter variation, subject of course to (2) being satisfied.

The resource underutilization of Example 1 may be further understood as resulting from the nonsmooth nature of the call arrival process. Consider the analogous *fluid model* in which ‘work’ of each type r arrives steadily at rate κ_r and is processed at rate $b_r(\mathbf{n})$, where each n_r is now the volume of work of type r in the network and where \mathbf{b} is again subject to constraints of the form of (1). Then, under condition (3) of Example 1, it is easy to see that every Pareto efficient control is stable, in the sense here that the total volume of work in the system eventually becomes and remains 0. For our stochastic model, the possible modified control discussed above, in which $b_1(\mathbf{n})$ is kept small whenever n_1 is small and n_2 is large, may be seen as a smoothing operation forcing the behaviour of the stochastic model to follow more closely that of the fluid model. However, as Example 2 below shows, instability may also occur in ways such that even the corresponding control for the analogous fluid model is also unstable.

For any control \mathbf{b} and any function f on $\mathbb{Z}_+^{|R|}$, define the function $D_{\mathbf{b}}f$ on $\mathbb{Z}_+^{|R|}$ by

$$D_{\mathbf{b}}f(\mathbf{n}) = \sum_{r \in R} [v_r(f(\mathbf{n} + \mathbf{e}_r) - f(\mathbf{n})) + \mu_r b_r(\mathbf{n})(f(\mathbf{n} - \mathbf{e}_r) - f(\mathbf{n}))]. \tag{5}$$

(Since, for \mathbf{n} and r such that $n_r = 0$, we also have $b_r(\mathbf{n}) = 0$, there is no problem arising from the lack of a formal definition of $f(\mathbf{n} - \mathbf{e}_r)$ in this case. Furthermore, $D_{\mathbf{b}}$ may be thought of as the generator of the Markov process $\mathbf{n}(\cdot)$ under the control \mathbf{b} .)

Example 2. Consider the network defined by $R = \{1, 2, 3\}$, $J = \{1, 2, 3\}$, and

$$A = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix},$$

with $\nu_r = \nu$, $\mu_r = 1$ for all r , and $c_j = c$ for all j . Assume that condition (2) is satisfied, i.e. that $2\nu < c$. Consider any control \mathbf{b} such that, for all $\mathbf{n} \neq \mathbf{0}$, we have $b_r(\mathbf{n}) = c$ for some r such that $n_r > 0$ and $b_s(\mathbf{n}) = 0$ for $s \neq r$. Then \mathbf{b} is Pareto efficient. However, for the function $f(\mathbf{n}) := \sum_{r \in R} n_r$ and for all $\mathbf{n} \neq \mathbf{0}$, we have $D_{\mathbf{b}}f(\mathbf{n}) = 3\nu - c$. Standard martingale arguments—again, see Asmussen (2003, Chapter 1, Proposition 5.4)—now show that \mathbf{b} is unstable whenever $3\nu \geq c$. In this example the instability is not simply the result of poor control for \mathbf{n} close to the boundary of \mathbb{Z}_+^3 , and is equally present in the analogous fluid model. We return to this example in Section 3.

Our aim in the remainder of the paper is to identify more general classes of stable controls, to provide insight into how the phenomenon of entrainment rises, and to show how controls may be modified if necessary so as to avoid it. In Section 2 we consider a number of fairly simple network topologies, and use Lyapunov function techniques to show that, for these, any Pareto efficient control \mathbf{b} is stable, provided *only* that, as discussed above and for suitable r , $b_r(\mathbf{n})$ is modified so as to be small whenever n_r is small. In Section 3 we consider controls which possess a natural monotonicity property, likely to be satisfied in any applications. We introduce some new analytical techniques to prove some fairly general results for the stability of the associated Markov chains. These give sufficient conditions for stability, which, for many classes of control, are also close to being necessary. These results are applicable to a substantial class of priority-based controls.

2. Simple network topologies

In this section we consider a number of network topologies and in each case show that, under condition (2), *any* Pareto efficient control is stable provided only that it is suitably modified for values of \mathbf{n} close to the boundary of $\mathbb{Z}_+^{|\mathbf{R}|}$.

The results of this section are based on simple Lyapunov function techniques, in particular Foster’s criterion. Proposition 1, below, states the specialisation of this to the present problem; see, for example, Asmussen (2003, Chapter 1, Proposition 5.3(ii)), noting that here jumps of the process $\mathbf{n}(\cdot)$ may only occur between neighbouring states, and also that the usual uniformization argument translates statements for discrete-time processes to the present continuous-time setting.

Proposition 1. *Given any control \mathbf{b} , suppose that there exists a positive function f on $\mathbb{Z}_+^{|\mathbf{R}|}$, a finite subset F of $\mathbb{Z}_+^{|\mathbf{R}|}$, and some $\varepsilon > 0$ such that $D_{\mathbf{b}}f(\mathbf{n}) \leq -\varepsilon$ for all $\mathbf{n} \notin F$. Then \mathbf{b} is stable.*

In considering the various network topologies of this section, it is convenient to define, for any $a \geq 1$, the function g_a on \mathbb{Z}_+ by

$$g_a(n) = \begin{cases} \frac{a}{2} + \frac{n^2}{2a} & \text{if } n < a, \\ n & \text{if } n \geq a. \end{cases}$$

We first consider again the very simple Example 1 of the Introduction, but with a general control \mathbf{b} .

Example 3. Let R , J , and A be as defined in Example 1 and again assume that $c_2 \leq c_1$ and also assume condition (3). We show that, given any $\delta > 0$, there exists some $a \geq 1$ (depending on the parameters ν_r , μ_r , c_j , and δ) such that a sufficient condition for the stability of any Pareto efficient control \mathbf{b} is given by

$$\kappa_2 - b_2(\mathbf{n}) \leq -\delta \quad \text{whenever } n_1 < a \text{ and } \mathbf{n} \notin F, \tag{6}$$

for some finite set F (necessarily including $\{\mathbf{n} : n_1 < a, n_2 = 0\}$). (Note that, from (3), condition (6) may be satisfied for any $\delta \in (0, c_2 - \kappa_2]$.)

Observe first that, from (3) and the saturation of resource 1 whenever $\mathbf{n} \neq \mathbf{0}$, there exists a $\delta' > 0$ such that, for any Pareto efficient control \mathbf{b} ,

$$\sum_{r=1}^2 (\kappa_r - b_r(\mathbf{n})) \leq -\delta' \quad \text{for all } \mathbf{n} \neq \mathbf{0}. \tag{7}$$

For any fixed a , define the function f_a on \mathbb{Z}_+^2 by

$$f_a(\mathbf{n}) = \frac{g_a(n_1)}{\mu_1} + \frac{n_2}{\mu_2}.$$

Elementary calculations show that, from (5), for any Pareto efficient control \mathbf{b} satisfying condition (6) for some finite F , and for any $\mathbf{n} \notin F$,

$$\begin{aligned} D_{\mathbf{b}} f_a(\mathbf{n}) &= \min\left(\frac{n_1}{a}, 1\right) [\kappa_1 - b_1(\mathbf{n})] + \kappa_2 - b_2(\mathbf{n}) + \frac{1}{2a} h_{1,a}(\mathbf{n}) \\ &\leq -\min\left(\frac{n_1}{a}, 1\right) \delta' - \max\left(1 - \frac{n_1}{a}, 0\right) \delta + \frac{1}{2a} h_{1,a}(\mathbf{n}) \end{aligned} \tag{8}$$

$$\leq -\min(\delta, \delta') + \frac{1}{2a} h_{1,a}(\mathbf{n}), \tag{9}$$

where the function $h_{1,a}$ is given by

$$h_{1,a}(\mathbf{n}) = \begin{cases} \kappa_1 + b_1(\mathbf{n}) & \text{if } 0 \leq n_1 < a, \\ b_1(\mathbf{n}) & \text{if } n_1 = a, \\ 0 & \text{if } n_1 > a, \end{cases} \tag{10}$$

and where (8) follows from (6) and (7). Since $b_1(\mathbf{n}) \leq c_1$ for any feasible control \mathbf{b} , it follows from (9) and (10) that a may be chosen sufficiently large that, for any Pareto efficient control \mathbf{b} satisfying (6) for some finite F ,

$$D_{\mathbf{b}} f_a(\mathbf{n}) \leq -\frac{1}{2} \min(\delta, \delta') \quad \text{for all } \mathbf{n} \notin F,$$

and so, by Proposition 1, \mathbf{b} is stable.

Thus, for this example, any Pareto efficient control is stable provided only that, outside of some finite set F , it is suitably modified for values of \mathbf{n} such that n_1 is small. However, in practice the parameters ν_r in particular are unlikely to be known, and so it is desirable to choose controls whose stability is robust. One such possibility is to choose any Pareto efficient control which assigns complete priority to calls of type 2 whenever $n_1 < a$ for some a . Provided

that only condition (3) is satisfied, the above result shows that this will be stable provided a is sufficiently large. The precise value of a required depends on the slack in the inequalities in (3), but calculations for the ‘worst case’, in which calls of type 1 have complete priority whenever $n_1 \geq a$, show that in general a need only be small.

Note, in particular, that the various fair-sharing controls defined in the introduction always satisfy (6) for some δ and for some sufficiently large F , and hence (as already remarked) are always stable provided that (3) is satisfied.

We now extend the above example to each of two more general network topologies. In each case it is again the case that only small modifications, identified below, are required to Pareto efficient controls in order to ensure their stability.

Example 4. Consider the network with $R = \{1, \dots, |R|\}$, and in which each call type r requires service from a dedicated resource of capacity c_r together with service from a resource which is shared by all call types and has capacity c_0 . The constraints, (1), defining the feasible controls \mathbf{b} are thus

$$\sum_{r \in R} b_r(\mathbf{n}) \leq c_0, \quad b_r(\mathbf{n}) \leq c_r, \quad r \in R.$$

We assume, without loss of generality, that $c_0 < \infty$ and that

$$c_0 \leq \sum_{r \in R} c_r. \tag{11}$$

As usual we assume condition (2), which here becomes

$$\sum_{r \in R} \kappa_r < c_0, \tag{12}$$

$$\kappa_r < c_r, \quad r \in R. \tag{13}$$

The case in which $|R| = 2$ with $c_1 = \infty$ is the earlier Example 3 (with c_0 here corresponding to c_1 there). However, in the general case described above a little more care is needed in the conditions for the stability of Pareto efficient controls.

For any $\mathbf{n} \in \mathbb{Z}_+^{|R|}$, define $n_{\min} = \min_{r \in R} n_r$. We briefly show that, given any $\delta > 0$, there again exists some $a \geq 1$ (depending on the parameters ν_r, μ_r, c_j , and δ) such that a sufficient condition for the stability of any Pareto efficient control \mathbf{b} is given by

$$\sum_{r: n_r \geq a'} (\kappa_r - b_r(\mathbf{n})) \leq -\delta \quad \text{for all } a' \in [1, a] \text{ and } \mathbf{n} \text{ such that } n_{\min} < a, \mathbf{n} \notin F, \tag{14}$$

for some finite set F . (The existence of a $\delta > 0$ such that condition (14) may be satisfied is guaranteed by (12) and (13). Furthermore, it is not difficult to see that for the case in which $R = 2$, this condition reduces to that of Example 3 taken together with a similar condition in which the roles of calls of types 1 and 2 are interchanged.) The proof of this result is similar to that of Example 3. First note that, from (11), it follows that, for any Pareto efficient control, the resource 0 is necessarily saturated for any \mathbf{n} such that $n_{\min} \geq 1$. Hence, from (13), there exists a $\delta' > 0$ such that, again, for any Pareto efficient control \mathbf{b} ,

$$\sum_{r \in R} (\kappa_r - b_r(\mathbf{n})) \leq -\delta' \quad \text{for all } \mathbf{n} \text{ such that } n_{\min} \geq 1. \tag{15}$$

For any fixed a , define the function f_a on $\mathbb{Z}_+^{|R|}$ by

$$f_a(\mathbf{n}) = \sum_{r \in R} \frac{g_a(n_r)}{\mu_r}.$$

Elementary calculations show that, from (5), for any Pareto efficient control \mathbf{b} satisfying (14) for some finite F , and for any $\mathbf{n} \notin F$,

$$\begin{aligned} D_{\mathbf{b}} f_a(\mathbf{n}) &= \sum_{r \in R} \min\left(\frac{n_r}{a}, 1\right) [\kappa_r - b_r(\mathbf{n})] + \frac{1}{2a} \sum_{r \in R} h_{r,a}(n_r) \\ &= \frac{1}{a} \sum_{a'=1}^a \sum_{r: n_r \geq a'} (\kappa_r - b_r(\mathbf{n})) + \frac{1}{2a} \sum_{r \in R} h_{r,a}(n_r) \\ &\leq -\min(\delta, \delta') + \frac{1}{2a} \sum_{r \in R} h_{r,a}(n_r), \end{aligned} \tag{16}$$

where, for each r , the function $h_{r,a}$ is given by (10) with the index r replacing the index 1, and where (16) follows from (14) for \mathbf{n} such that $n_{\min} < a$ and from (15) for \mathbf{n} such that $n_{\min} \geq a$. It now follows, as in Example 3, that a may be chosen sufficiently large that, for any Pareto efficient control \mathbf{b} satisfying (14) for some finite F ,

$$D_{\mathbf{b}} f_a(\mathbf{n}) \leq -\frac{1}{2} \min(\delta, \delta') \quad \text{for all } \mathbf{n} \notin F,$$

and so, again by Proposition 1, the control \mathbf{b} is stable.

Thus, in order to ensure the stability of a general Pareto efficient control \mathbf{b} , it is only necessary to appropriately modify $\mathbf{b}(\mathbf{n})$ for \mathbf{n} such that $n_{\min} < a$. In particular, it follows from the above result that a Pareto efficient control \mathbf{b} whose stability is reasonably robust is given by requiring that, for some a , for all \mathbf{n} (such that $n_{\min} < a$), and for all $a' \in [1, a]$, calls of type r , such that $n_r \geq a'$, collectively have complete priority over calls of the remaining types, that is, requiring that

$$\sum_{r: n_r \geq a'} b_r(\mathbf{n}) = \min\left(c_0, \sum_{r: n_r \geq a'} c_r\right).$$

For sufficiently large a , depending on the slack in the inequalities in (12) and (13), any such control is stable.

Now consider further the case in which $R = 2$. Note that the present topology is completely general for a network with two call types. Let \mathbf{b} be any Pareto efficient control such that

$$\lim_{n_1 \rightarrow \infty} b_1(n_1, n_2) = c_1 \quad \text{for all } n_2, \tag{17}$$

$$\lim_{n_2 \rightarrow \infty} b_2(n_1, n_2) = c_2 \quad \text{for all } n_1. \tag{18}$$

Then it is straightforward to show that, for $\delta < \min_{r=1,2}(c_r - \kappa_r)$ and for any a , (14) is satisfied for sufficiently large F , and so \mathbf{b} is stable.

In particular, (17) and (18) are again satisfied by the various fair-sharing controls. The present conditions are of course considerably more general. However, the development of corresponding results for general networks with $R \geq 3$ remains a challenging problem.

Our final network topology is a simple ‘backbone’ structure.

Example 5. Consider a network with resource set $J = \{1, \dots, k\}$, where resource j has capacity c_j as usual. The set of call types is given by $R = \{0, 1, \dots, k\}$, where calls of each type $r = 1, \dots, k$ require service from the single resource $j = r$, while calls of type 0 require service from each of the resources $1, \dots, k$.

The constraints, (1), defining the feasible controls \mathbf{b} are given by

$$b_0(\mathbf{n}) + b_r(\mathbf{n}) \leq c_r, \quad r = 1, \dots, k.$$

Again we assume condition (2), which here becomes

$$\kappa_0 + \kappa_r < c_r, \quad r = 1, \dots, k. \tag{19}$$

The state of the network is thus denoted here by $\mathbf{n} = (n_0, n_1, \dots, n_k)$. For any such \mathbf{n} , define $\hat{n}_{\max} = \max(n_1, \dots, n_k)$.

This example again generalises that of Example 3 with the call type 0 here playing the role of the call type 2 in that example. In a fairly straightforward generalisation of the result of that example, it is here the case that, given any $\delta > 0$, there again exists an $a \geq 1$ such that a sufficient condition for the stability of any Pareto efficient control \mathbf{b} is given by

$$\kappa_0 - b_0(\mathbf{n}) \leq -\delta \quad \text{whenever } \hat{n}_{\max} < a \text{ and } \mathbf{n} \notin F, \tag{20}$$

for some finite set F . To see this, first suppose that $\mu_r = 1$ for all $r = 0, 1, \dots, k$ and that (20) is satisfied. For any $a \geq 1$, define the function f_a on $\mathbb{Z}_+^{|R|}$ by $f_a(\mathbf{n}) = n_0 + g_a(\hat{n}_{\max})$. Then, as in the case of Example 3, it follows straightforwardly from (19) and (20) that there exists a $\delta' > 0$ and a sufficiently large a that $D_{\mathbf{b}} f_a(\mathbf{n}) \leq -\delta'$ for all $\mathbf{n} \in \mathbb{Z}_+^{|R|} \setminus F$ such that additionally $\hat{n}_{\max} = n_r$ for a *single* value of $r = 1, \dots, k$. It is now easy to see that this latter restriction may be removed (possibly at the expense of increasing a) by suitably smoothing the function f in the neighbourhood of those \mathbf{n} such that $\hat{n}_{\max} = n_r$ for two or more values of $r = 1, \dots, k$. The desired result thus follows in this case; for general μ_r only routine modifications to the above argument are required.

Hence, we again have that any Pareto efficient control requires only slight modification—for those \mathbf{n} such that $\hat{n}_{\max} < a$ —in order to be stable. A robust Pareto efficient control is given, for example, by assigning complete priority to calls of type 0 whenever $\hat{n}_{\max} < a$, the necessary value of a depending on the slack in the inequalities (19).

Examples 3–5, above, all consider fairly simple network topologies. In the analogous fluid model defined in the introduction, it is easily seen that, for each of these topologies, condition (2) is sufficient for the stability of *any* Pareto efficient control. (In each case this follows, for example, by using the same Lyapunov function as for the stochastic model, except that the function g_a may be replaced by the identity function.) The examples illustrate a principle which seems likely to be true for more general network topologies, namely that when a control is such that it is stable for the fluid model, then there is a closely approximating control which is stable for the corresponding stochastic model.

In the next section we consider stability criteria for quite general network topologies, applicable typically to controls where there is some prioritisation among call types.

3. Stability of monotonic controls

Many controls likely to be of practical application possess a simple monotonicity property (see below). In this section we study stability for a wide class of such controls, giving sufficient

conditions for stability, which, for many classes of control, are also close to being necessary (see Remark 2, below, which further discusses the applicability of the results of this section). We require first the following quite general lemma.

We shall say that a control \mathbf{b} is *bounded* if, for all r , $b_r(\mathbf{n})$ is bounded in \mathbf{n} .

Lemma 1. *Let \mathbf{b} be any bounded control and, as usual, let $\mathbf{n}(\cdot)$ denote the Markov process for the corresponding network. Then, for all $r \in R$,*

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t b_r(\mathbf{n}(u)) \, du \leq \kappa_r \quad \text{almost surely (a.s.).} \tag{21}$$

Proof. For all $r \in R$, the compensated process $n_r^*(\cdot)$ defined by

$$n_r^*(t) = n_r(t) - n_r(0) + \int_0^t (b_r(\mathbf{n}(u)) - \kappa_r) \, du$$

is a zero-mean martingale. Also, since the transition rates of the Markov process $\mathbf{n}(\cdot)$ are bounded, it follows that, for some constant M and for all $t \geq 0$, we have $E(n_r^*(t)^2) \leq Mt$. Thus, for $\frac{1}{2} < \alpha < 1$, the process $(n_r^*(t)/t^\alpha)_{t>0}$ is an \mathcal{L}^2 -bounded supermartingale, and so, as $t \rightarrow \infty$, converges almost surely to some finite random variable. It follows that

$$\frac{n_r(t)}{t} + \frac{1}{t} \int_0^t b_r(\mathbf{n}(u)) \, du \rightarrow \kappa_r \quad \text{a.s.}$$

Since the process $n_r(\cdot)$ is positive, (21) now follows.

Remark 1. In the case where the bounded control \mathbf{b} is stable, the Markov process $\mathbf{n}(\cdot)$ is positive recurrent, and so we have the stronger result that, for all $r \in R$,

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t b_r(\mathbf{n}(u)) \, du = E_\pi b_r = \kappa_r \quad \text{a.s.} \tag{22}$$

Here $E_\pi b_r$ denotes the expectation of the function b_r with respect to the stationary distribution π of $\mathbf{n}(\cdot)$. The first equality in (22) follows from the ergodic theorem, and the second is simply the assertion that, under stationarity (and readily deducible from the balance equations defining π), the expected arrival and departure rates for calls of type r are equal.

We shall say that a bounded control \mathbf{b} is *monotonic* if, for all r and for all \mathbf{n} ,

$$\begin{aligned} b_r(\mathbf{n}) &\text{ is increasing in } n_r \text{ (with } n_s \text{ fixed for all } s \neq r), \\ b_r(\mathbf{n}) &\text{ is decreasing in } n_{s'} \text{ (with } n_{s'} \text{ fixed for all } s' \neq s) \text{ for all } s \neq r. \end{aligned}$$

Note that, depending on the network structure, this property is natural in many applications. For instance, for the structure of Example 4, it is possessed by all the fair-sharing controls, and also by any other reasonable Pareto efficient control. For more complex network structures, controls may be coupled to monotonic controls to establish stability results using the results given below. See also Example 6, below. We further note that a related but somewhat different definition of monotonicity is used by Bonald and Proutière (2004).

For any monotonic control \mathbf{b} , and for each $S \subseteq R$ (including the case where S is the empty set \emptyset), define the function $\mathbf{b}^S: \mathbb{Z}_+^{|S|} \rightarrow \mathbb{R}_+^{|R|}$ by

$$b_r^S(\mathbf{n}_S) = \lim_{n_s \rightarrow \infty \text{ for all } s \notin S} b_r(\mathbf{n}), \quad r \in S, \tag{23}$$

$$b_r^S(\mathbf{n}_S) = \lim_{n_r \rightarrow \infty} \lim_{n_s \rightarrow \infty \text{ for all } s \notin S \cup \{r\}} b_r(\mathbf{n}), \quad r \notin S, \tag{24}$$

where $\mathbf{n}_S = (n_s, s \in S)$ and where, in taking the limits on the right-hand side of (23) and (24), the vector of those coordinates of \mathbf{n} that belong to S is held fixed at \mathbf{n}_S . Note that, by monotonicity, this function is well-defined. In particular, in (23), the order within $R \setminus S$ in which the limits are taken is irrelevant; however, in (24), the final limit to be taken must be $n_r \rightarrow \infty$.

For any monotonic control \mathbf{b} and any S as above, we shall say that \mathbf{b}^S is *stable* if the application of the control $(b_s^S, s \in S)$ to calls in the network of type $s \in S$ (with ν_s, μ_s , and $s \in S$, as usual) yields a positive recurrent Markov process $\mathbf{n}^S(\cdot) = (n_s^S(\cdot), s \in S)$. (This $|S|$ -dimensional process may be thought of as that which results when the number of calls of each type $r \notin S$ is infinite.) When \mathbf{b}^S is stable we shall let π^S denote the stationary distribution of $\mathbf{n}^S(\cdot)$ (or, where necessary, the probability function of this distribution); we shall further define, for each $r \in R$,

$$E_{\pi^S} b_r^S = \sum_{\mathbf{n}_S \in \mathbb{Z}_+^{|S|}} \pi^S(\mathbf{n}_S) b_r^S(\mathbf{n}_S) \tag{25}$$

to be the expected value of b_r^S under this distribution. In the case where S is the empty set \emptyset , we have $\mathbf{b}^\emptyset = (b_r^\emptyset, r \in R)$ is a vector of constants. We make the natural convention that \mathbf{b}^\emptyset is always stable; the distribution π^\emptyset is concentrated on a single point, and we have $E_{\pi^\emptyset} b_r^\emptyset = b_r^\emptyset$ for all $r \in R$.

Our main result of this section is Theorem 1, below. The first part is similar in spirit to the results of Borovkov (1998, Chapter 8) for asymptotically spatially homogeneous Markov processes. However, the application of those results here would require that the right-hand side of (24) is invariant under interchange of the limits in that expression, a condition which is not in general satisfied for our monotonic controls. Rather the monotonicity itself provides sufficient structure to obtain the results of Theorem 1.

Theorem 1. *Suppose that the control \mathbf{b} is monotonic and that $S \subseteq R$ is such that \mathbf{b}^S is stable. We have the following two cases.*

(i) *If $r \notin S$ is such that*

$$E_{\pi^S} b_r^S > \kappa_r, \tag{26}$$

then $\mathbf{b}^{S \cup \{r\}}$ is stable.

(ii) *If $r \notin S$ is such that*

$$E_{\pi^S} b_r^S < \kappa_r, \tag{27}$$

then $\mathbf{b}^{S \cup \{r\}}$ is unstable.

Remark 2. Given the stability of \mathbf{b}^S for some $S \subset R$ (recall that, as already remarked, \mathbf{b}^\emptyset is always stable), Theorem 1 gives criteria for determining the stability or otherwise of $\mathbf{b}^{S \cup \{r\}}$ for any $r \notin S$, except only in the case of equality in (26) or (27) (where the natural conjecture is that $\mathbf{b}^{S \cup \{r\}}$ is unstable; see also the remarks at the end of Example 6). Recursive application of the theorem thus yields sufficient conditions both for the stability and the instability of monotonic controls. However, note that, for example, for the case in which $R = \{1, 2\}$, $\mathbf{b}^{\{1\}}$ and $\mathbf{b}^{\{2\}}$ may both be unstable, while $\mathbf{b} = \mathbf{b}^{\{1,2\}}$ is stable, as is the case for fair-sharing controls here. In such circumstances Theorem 1 does not settle the question of the stability of \mathbf{b} . Rather its primary application is to controls in which there is a sufficient hierarchy of prioritisation as to permit the recursive application of the first part of the theorem to at least establish the stability of \mathbf{b}^S for $S = R \setminus \{r\}$ for some $r \in R$. The theorem then also (in general) settles the question of the stability of \mathbf{b} itself. For an illustration of the application of the theorem see Example 6.

Proof of Theorem 1. Since, for given $S \subset R$ and $r \notin S$, the stability of $\mathbf{b}^{S \cup \{r\}}$ corresponds to the positive recurrence of the Markov process $\mathbf{n}^{S \cup \{r\}}(\cdot)$ defined above (in which the number of calls of each type $s \notin S \cup \{r\}$ is effectively held at ∞), it is sufficient to prove (i) and (ii) for the case in which $S = R \setminus \{r'\}$ for some r' , and we henceforth assume this. (The primary advantage of doing so is that we avoid some unpleasant notational complexity.) We identify any $\mathbf{n} \in \mathbb{Z}_+^{|R|}$ with the pair $(\mathbf{n}_S, n_{r'})$, where $\mathbf{n}_S = (n_s, s \in S)$. Recall that, for each such \mathbf{n}_S , we then have $b_{r'}^S(\mathbf{n}_S) = \lim_{n_{r'} \rightarrow \infty} b_{r'}(\mathbf{n}_S, n_{r'})$.

First suppose that (26) holds. We need to show that \mathbf{b} is stable. The underlying idea is that the monotonicity of \mathbf{b} and stability of \mathbf{b}^S ensure that the components $(n_s(\cdot), s \in S)$ of the process $\mathbf{n}(\cdot)$ become and remain small, and that (26) then ensures that, except in some finite region A , the remaining component $n_{r'}(\cdot)$ of this process is decreasing at a rate which is bounded away from 0; thus, the process $\mathbf{n}(\cdot)$ spends, in the long term, a nonzero proportion of time within A .

It follows from (26) and the monotonicity of \mathbf{b} that we can choose a finite set $A = \{\mathbf{n} : n_r \leq \bar{n}_r, r \in R\} \subset \mathbb{Z}_+^{|R|}$ and a positive function $\bar{b}_{r'}$ on $\mathbb{Z}_+^{|S|}$ such that

$$\bar{b}_{r'} \text{ is decreasing in each of its arguments,} \tag{28}$$

$$\sum_{\mathbf{n}_S \in \mathbb{Z}_+^{R-1}} \pi^S(\mathbf{n}_S) \bar{b}_{r'}(\mathbf{n}_S) > \kappa_{r'}, \tag{29}$$

$$b_{r'}(\mathbf{n}) \geq \bar{b}_{r'}(\mathbf{n}_S) \text{ for all } \mathbf{n} \notin A. \tag{30}$$

For example, given $\bar{n}_r, r \in R$, we may choose

$$\bar{b}_{r'}(\mathbf{n}_S) = \mathbf{1}(n_s \leq \bar{n}_s, s \in S) b_{r'}(\mathbf{n}_S, \bar{n}_{r'}), \quad \mathbf{n}_S \in \mathbb{Z}_+^{|S|}, \tag{31}$$

where $\mathbf{1}(\cdot)$ denotes the indicator function. Condition (28) then follows from the monotonicity of \mathbf{b} ; also, since \mathbf{b} is nonnegative, (30) follows trivially from (31), except for \mathbf{n} such that $n_s \leq \bar{n}_s$ for $s \in S$ and $n_{r'} > \bar{n}_{r'}$, in which case (30) again follows from the monotonicity of \mathbf{b} ; finally, (29) follows from (26) and the monotone convergence theorem by choosing $\bar{n}_r, r \in R$, all sufficiently large, since, for any \mathbf{n}_S ,

$$\lim_{\bar{n}_r \rightarrow \infty} \mathbf{1}(n_s \leq \bar{n}_s, s \in S) b_{r'}(\mathbf{n}_S, \bar{n}_{r'}) = \lim_{\bar{n}_{r'} \rightarrow \infty} b_{r'}(\mathbf{n}_S, \bar{n}_{r'}) = b_{r'}^S(\mathbf{n}_S).$$

It also follows from the monotonicity of the control \mathbf{b} that we can couple the corresponding Markov process $\mathbf{n}(\cdot)$ to a process $\mathbf{n}^S(\cdot) = (n_s^S(\cdot), s \in S)$ on $\mathbb{Z}_+^{|S|}$ with control \mathbf{b}^S in such a way that $n_s(t) \leq n_s^S(t)$ for all $t > 0$ and for all $s \in S$. Also, since the process $\mathbf{n}^S(\cdot)$ has stationary distribution π^S , it follows from (28) that, for the function $\bar{b}_{r'}$ defined above,

$$\limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t \bar{b}_{r'}(\mathbf{n}_S(u)) \, du \geq \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \bar{b}_{r'}(\mathbf{n}^S(u)) \, du = E_{\pi^S} \bar{b}_{r'}, \tag{32}$$

where $\mathbf{n}_S(\cdot) = (n_s(\cdot), s \in S)$ and where the final equality follows from the ergodic theorem.

It now follows from Lemma 1 that, for some $M > 0$,

$$\begin{aligned} \kappa_{r'} &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t b_{r'}(\mathbf{n}(u)) \, du \\ &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t b_{r'}(\mathbf{n}(u)) \mathbf{1}(\mathbf{n}(u) \notin A) \, du \\ &\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t \bar{b}_{r'}(\mathbf{n}_S(u)) \mathbf{1}(\mathbf{n}(u) \notin A) \, du \end{aligned} \tag{33}$$

$$\geq \limsup_{t \rightarrow \infty} \frac{1}{t} \int_0^t \bar{b}_{r'}(\mathbf{n}_S(u)) \, du - \lim_{t \rightarrow \infty} \frac{M}{t} \int_0^t \mathbf{1}(\mathbf{n}(u) \in A) \, du \tag{34}$$

$$\geq E_{\pi^S} \bar{b}_{r'} - \lim_{t \rightarrow \infty} \frac{M}{t} \int_0^t \mathbf{1}(\mathbf{n}(u) \in A) \, du, \tag{35}$$

where the inequality in (33) follows from (30), the inequality in (34) follows from (28) since the function $\bar{b}_{r'}$ is necessarily bounded (note that, since A is finite, the limit in the final term in (34) always exists), and the final inequality in (35) follows from (32). Thus, from (35) and (29), we have

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbf{1}(\mathbf{n}(u) \in A) \, du > 0.$$

Since A is finite it now follows from the ergodic theorem that the Markov process $\mathbf{n}(\cdot)$ is positive recurrent and so \mathbf{b} is stable.

Now suppose, instead, that (27) holds. We show that the Markov process $\mathbf{n}(\cdot)$ corresponding to \mathbf{b} is transient (and hence \mathbf{b} is unstable). The underlying idea here is that whenever $n_{r'}(\cdot)$ is very large, the process $\mathbf{n}(\cdot)$ again behaves approximately as if it were controlled by \mathbf{b}^S , and thus, from (27), we may expect that $\lim_{t \rightarrow \infty} n_{r'}(t) = \infty$ a.s. To make this rigorous we again use the monotonicity of \mathbf{b} to couple the process $\mathbf{n}(\cdot)$ to a process $\hat{\mathbf{n}}(\cdot)$ whose control is sufficiently close to that of \mathbf{b}^S that we may show that $\lim_{t \rightarrow \infty} \hat{n}_{r'}(t) = \infty$ a.s., and for which the coupling also ensures that $\lim_{t \rightarrow \infty} n_{r'}(t) = \infty$ with strictly positive probability.

Given $\bar{n}_{r'} \in \mathbb{Z}_+$ (fixed, to be chosen later), define a Markov process $\hat{\mathbf{n}}(\cdot) = (\hat{n}_r(\cdot), r \in R)$ as follows: for each $s \in S$, the component process $\hat{n}_s(\cdot)$ has state space \mathbb{Z}_+ as usual, while $\hat{n}_{r'}(\cdot)$ has state space \mathbb{Z} ; for each $r \in R$, transitions $n_r \rightarrow n_r + 1$ occur at rate ν_r as usual, while transitions $n_r \rightarrow n_r - 1$ occur at rate $\mu_r \hat{b}_r(\mathbf{n})$, where

$$\begin{aligned} \hat{b}_s(\mathbf{n}) &= b_s(\mathbf{n}_S, \bar{n}_{r'}), & s \in S, \\ \hat{b}_{r'}(\mathbf{n}) &= b_{r'}^S(\mathbf{n}_S). \end{aligned}$$

Observe that the process $\hat{\mathbf{n}}(\cdot)$ has uniformly bounded transition rates which are independent of $n_{r'} \in \mathbb{Z}$. Suppose that

$$\mathbf{n}(0) = \hat{\mathbf{n}}(0), \quad n_{r'}(0) > \bar{n}_{r'}. \tag{36}$$

Define the random time $T = \min\{t > 0 : n_{r'}(t) < \bar{n}_{r'}\}$. It follows from the monotonicity of \mathbf{b} that, for all \mathbf{n} such that $n_{r'} \geq \bar{n}_{r'}$,

$$\hat{b}_s(\mathbf{n}) \geq b_s(\mathbf{n}) \geq b_s^S(\mathbf{n}_S), \quad s \in S, \tag{37}$$

and, hence, that we may couple the processes $\hat{\mathbf{n}}(\cdot)$, $\mathbf{n}(\cdot)$, and the $|S|$ -dimensional process $\mathbf{n}^S(\cdot)$ with control $(b_s^S, s \in S)$ in such a way that

$$\hat{n}_s(t) \leq n_s^S(t) \quad \text{for all } s \in S \text{ and for all } t \geq 0 \tag{38}$$

and

$$\hat{n}_s(t) \leq n_s(t) \quad \text{for all } s \in S \text{ and for all } 0 \leq t \leq T. \tag{39}$$

Since $\mathbf{n}^S(\cdot)$ is assumed positive recurrent with stationary distribution π^S , it follows, from (37) or (38), that $\hat{\mathbf{n}}_S(\cdot) = (\hat{n}_s(\cdot), s \in S)$ is similarly positive recurrent with stationary distribution $\hat{\pi}_S$ say. Furthermore, as $\bar{n}_{r'} \rightarrow \infty$, the control $\hat{\mathbf{b}}$ converges pointwise in each of its components to \mathbf{b}^S . Hence, elementary arguments (e.g. consideration of the times of return to $\mathbf{0}$ of the

process $\mathbf{n}^S(\cdot)$, coupled with the use of (38) and the ergodic theorem) show that, again as $\bar{n}_{r'} \rightarrow \infty$, $\hat{\pi}_S$ converges in distribution to π^S . Also, since $b_{r'}^S$ is bounded, it now follows, from (27), that we may choose the constant $\bar{n}_{r'}$ sufficiently large that

$$E_{\hat{\pi}_S} b_{r'}^S < \kappa_{r'} \tag{40}$$

(where, analogously to (25), $E_{\hat{\pi}_S} b_{r'}^S$ is the expectation of $b_{r'}^S$ with respect to $\hat{\pi}_S$).

It further follows, from the monotonicity of \mathbf{b} and from (39), that, for all $0 \leq t \leq T$,

$$\begin{aligned} \hat{b}_{r'}(\hat{\mathbf{n}}(t)) &= b_{r'}^S(\hat{\mathbf{n}}_S(t)) \\ &\geq b_{r'}^S(\mathbf{n}_S(t)) \end{aligned} \tag{41}$$

$$\geq b_{r'}(\mathbf{n}(t)), \tag{42}$$

where (41) follows from (39) and the monotonicity of \mathbf{b} while (42) follows from the definition of \mathbf{b}^S and, again, the monotonicity of \mathbf{b} . It now follows, from (42), that we may also couple the components r' of the processes $\hat{\mathbf{n}}(\cdot)$ and $\mathbf{n}(\cdot)$ in such a way that

$$\hat{n}_{r'}(t) \leq n_{r'}(t) \quad \text{for all } 0 \leq t \leq T. \tag{43}$$

As noted above, the process $\hat{\mathbf{n}}_S(\cdot) = (\hat{n}_s(\cdot), s \in S)$ has stationary distribution $\hat{\pi}_S$, while the process $\hat{n}_{r'}(\cdot)$ may be viewed as a Markov additive process modulated by the remaining components $\hat{\mathbf{n}}_S(\cdot)$ of $\hat{\mathbf{n}}(\cdot)$. From (40), the expectation of the increments of $\hat{n}_{r'}(\cdot)$ between those times at which $\hat{\mathbf{n}}(\cdot)$ returns to any fixed state is strictly positive. It follows, from the standard theory of Markov additive processes, that $\lim_{t \rightarrow \infty} \hat{n}_{r'}(t) = \infty$ a.s., and furthermore that, under condition (36),

$$P(\hat{n}_{r'}(t) \geq \bar{n}_{r'} \text{ for all } t \geq 0, \lim_{t \rightarrow \infty} \hat{n}_{r'}(t) = \infty) > 0,$$

and, hence, from (43), it also follows that

$$P(n_{r'}(t) \geq \bar{n}_{r'} \text{ for all } t \geq 0, \lim_{t \rightarrow \infty} n_{r'}(t) = \infty) > 0.$$

Hence, the process $\mathbf{n}(\cdot)$ is transient as required.

We illustrate the use of the above result with a simple example.

Example 6. Consider again the network of Example 2, in which $R = J = 3$. As previously observed a necessary and sufficient condition for the existence of *some* stable control is given by $2\nu < c$. Furthermore, if $3\nu < c$, then Proposition 1 with the Lyapunov function f given by $f(\mathbf{n}) = \sum_{r=1}^3 n_r$ shows that *any* Pareto efficient control is stable. Suppose now that $2\nu < c$ and that the Pareto efficient control \mathbf{b} is such that, for $r = 1, 2$, $b_r(\mathbf{n})$ is independent of n_3 and

$$b_1(\mathbf{n}) + b_2(\mathbf{n}) = c \quad \text{for all } \mathbf{n} \text{ such that } \max(n_1, n_2) > 0. \tag{44}$$

Thus, in particular, calls of types 1 and 2 collectively have complete priority over calls of type 3. Although we do not in this example require any further monotonicity conditions on \mathbf{b} , it follows, from the requirement of Pareto efficiency, that the control $\mathbf{b}^{(1,2)}: \mathbb{Z}_+^2 \rightarrow \mathbb{R}_+^3$ is well defined, as before, being obtained from \mathbf{b} by letting $n_3 \rightarrow \infty$. It follows, from (44), that the condition $2\nu < c$ is necessary and sufficient to ensure that $\mathbf{b}^{(1,2)}$ is stable. We use (a slight modification of) Theorem 1 to investigate the stability of \mathbf{b} . The stationary distribution

$\pi^{\{1,2\}}$ on \mathbb{Z}_+^2 induced by $\mathbf{b}^{\{1,2\}}$ is here just that of the process $(n_1(\cdot), n_2(\cdot))$. Furthermore, since, from (44), $n_1(\cdot) + n_2(\cdot)$ is Markov, with a stationary distribution which is geometric and independent of any more detailed specification of $\mathbf{b}^{\{1,2\}}$, it follows that

$$\pi^{\{1,2\}}(0, 0) = 1 - \frac{2\nu}{c}.$$

It follows, from (1), (44), and the Pareto efficiency of \mathbf{b} , that

$$b_3^{\{1,2\}}(0, 0) = c, \quad b_3^{\{1,2\}}(n_1, 0) = 0 \quad \text{for all } n_1 \geq 1, \quad \text{and} \\ b_3^{\{1,2\}}(0, n_2) = 0 \quad \text{for all } n_2 \geq 1.$$

Thus, we have

$$E_{\pi^{\{1,2\}}} b_3^{\{1,2\}} = \sum_{(n_1, n_2) \in \mathbb{Z}_+^2} \pi^{\{1,2\}}(n_1, n_2) b_3^{\{1,2\}}(n_1, n_2) \geq c - 2\nu, \tag{45}$$

with equality if and only if $b_3^{\{1,2\}}(n_1, n_2) = 0$ for all (n_1, n_2) such that $\min(n_1, n_2) \geq 1$. But this latter condition holds if and only if, for all (n_1, n_2) such that $\min(n_1, n_2) \geq 1$, we have $\min(b_1^{\{1,2\}}(n_1, n_2), b_2^{\{1,2\}}(n_1, n_2)) = 0$, i.e. in the case of the control considered in Example 2 in which maximum resource is always allocated to calls of one type, and in which we have already observed that we have stability if and only if $3\nu < c$. Otherwise we have strict inequality in (45).

Now note that, although \mathbf{b} does not here satisfy all the conditions for monotonicity given earlier, the assumption that $b_1(\mathbf{n})$ and $b_2(\mathbf{n})$ are independent of n_3 ensures that Theorem 1 continues to apply, indeed in a slightly improved form, to show that the condition $E_{\pi^{\{1,2\}}} b_3^{\{1,2\}} > \nu$ is necessary and sufficient for the stability of \mathbf{b} . (For the sufficiency, note that the proof of part (i) of the theorem, with $S = \{1, 2\}$ and $r' = 3$, goes through as before, except that the coupling between $\mathbf{n}(\cdot)$ and $\mathbf{n}^S(\cdot)$ is now obtained with equality, and so we no longer require (28) in order to obtain (32). Similar obvious simplifications apply to the proof of part (ii), which here becomes a fairly standard argument and in particular delivers null recurrence—and hence instability—in the case $E_{\pi^{\{1,2\}}} b_3^{\{1,2\}} = \nu$.)

Suppose now that c and \mathbf{b} are held fixed and that ν is allowed to vary. The obvious coupling argument shows that if \mathbf{b} is stable for any ν then it is also stable for any $\nu' < \nu$. The above adaptation of Theorem 1, together with (45), shows that there is some critical parameter λ (depending on the detailed specification of $\mathbf{b}^{\{1,2\}}$ and hence $\pi^{\{1,2\}}$) such that $\frac{1}{3} \leq \lambda \leq \frac{1}{2}$ and that \mathbf{b} is stable if $\nu < \lambda c$ and unstable if $\nu > \lambda c$. For the control of Example 2, we already know that $\lambda = \frac{1}{3}$; otherwise for the case in which $\nu = c/3$ we have strict inequality in (45) and hence stability; simple continuity arguments now give $\lambda > \frac{1}{3}$ in this case.

Acknowledgement

The authors are most grateful to Serguei Foss and to Takis Konstantopoulos for some helpful discussions, and also to an anonymous referee for helpful comments and corrections.

References

ASMUSSEN, S. (2003). *Applied Probability and Queues*. Springer, New York.
 BONALD, T. AND MASSOULIÉ, L. (2001). Impact of fairness on Internet performance. In *Proc. ACM SIGMETRICS 2001*, ACM, New York, pp. 82–91.

- BONALD, T. AND PROUTIERE, A. (2004). On stochastic bounds for monotonic processor sharing networks. *Queueing Systems* **47**, 81–106.
- BOROVKOV, A. A. (1998). *Ergodicity and Stability of Stochastic Processes*. John Wiley, Chichester.
- DAI, J. (1995). On positive Harris recurrence of multiclass queueing networks: a unified approach via fluid limit models. *Ann. Appl. Prob.* **5**, 49–77.
- 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. AND WILLIAMS, R. J. (2004). Fluid model for bandwidth sharing. *Ann. Appl. Prob.* **14**, 1055–1083.
- KELLY, F. P., MAULLOO, A. K. AND TAN, D. K. H. (1998). Rate control in communication networks: shadow prices, proportional fairness and stability. *J. Operat. Res. Soc.* **49**, 237–252.
- MO, J. AND WALRAND, J. (2000). Fair end-to-end window-based congestion control. *IEEE/ACM Trans. Networking* **8**, 556–567.
- ROBERTS, J. W. AND MASSOULIÉ, L. (2000). Bandwidth sharing and admission control for elastic traffic. *Telecommun. Systems* **15**, 185–201.