11

The asymptotic behaviour of large loss networks

Stan Zachary

Heriot-Watt University

Abstract

We study the limit behaviour of controlled loss networks as capacity and offered traffic are allowed to increase in proportion, reviewing and extending recent work based on the functional law of large numbers of Hunt and Kurtz. We consider in detail single and two-resource networks.

1 Introduction

In this paper we study large loss networks in which the offered traffic is subject to acceptance controls. We review recent work of Hunt and Kurtz (1994), who established rigorous results for the asymptotic dynamics of such networks as capacity and offered traffic are allowed to increase in proportion, and we relate these results to asymptotic equilibrium behaviour. We further study the detailed behaviour of networks with at most two resources, extending results of Bean *et al.* (1994b, 1995), and giving some additional results.

The asymptotic results considered here have important applications to the control of modern communications networks, which are typically large and which may simultaneously carry traffic with very different capacity requirements and holding times. A failure to apply effective controls in such networks can lead to a serious degradation in performance.

The results also remain qualitatively correct for smaller capacity networks. Bean *et al.* (1994a, 1995) and Moretta (1995) derive refinements which permit more accurate modelling of the quantitative behaviour of networks of all capacities.

The mathematical framework is the same as that of Hunt and Kurtz (1994). Consider a sequence of loss networks, indexed by a scale parameter N. All members of the sequence are identical except in respect of capacities and call arrival rates (which, as defined more precisely by eqn (1.1) below, are essentially proportional to N), and are identically controlled. Resources (or links) are indexed in a finite set $\mathcal{J}$ and call types in a finite set $\mathcal{R}$. For the Nth member of the sequence, each resource $j \in \mathcal{J}$ has integer capacity $C_j(N)$, and calls of each type $r \in \mathcal{R}$ arrive as a Poisson process of rate $\kappa_r(N)$. Each such call, if accepted, simultaneously requires an integer A_{jr} units of the capacity of each resource j for the duration of its holding time, which is exponentially distributed with mean $1/\mu_r$. All arrival streams and holding times are independent.

Let $n^N(t) = (n_r^N(t),\ r \in \mathcal{R})$, where $n_r^N(t)$ is the number of calls of type r in progress at time t, and let $m^N(t) = (m_j^N(t),\ j \in \mathcal{J})$ where $m_j^N(t) = C_j(N) - \sum_{r\in\mathcal{R}} A_{jr} n_r^N(t)$ is the free capacity of resource j at time t. A call of type r arriving at time t is accepted if and only if $m^N(t-)$ belongs to some acceptance region $\mathcal{A}_r$, which we formally regard as a subset of the space $E = (\mathbb{Z}_+ \cup \{\infty\})^J$, where $J = |\mathcal{J}|$. (Of course, the process $m^N(\cdot)$ only takes values in $\mathbb{Z}_+^J$.) We further require that each set $\mathcal{A}_r$ is well-behaved in the sense that its indicator function $I_{\mathcal{A}_r}$ is continuous, where the topology of E is the product of the topology of the one-point compactification of $\mathbb{Z}_+$.

This framework permits the modelling of a wide variety of control mechanisms, including most of those, such as fixed routing, trunk reservation and alternative routing, employed in practical applications to communications networks. For details see Hunt and Kurtz (1994).

Suppose that, as $N \to \infty$, for all $j \in \mathcal{J}$, $r \in \mathcal{R}$,

$$\frac{1}{N} C_j(N) \to C_j, \qquad \frac{1}{N} \kappa_r(N) \to \kappa_r. \tag{1.1}$$

Then, under appropriate initial conditions, the normalized process $x^N(\cdot) = n^N(\cdot)/N$ might reasonably be expected to converge to a 'fluid limit' process $x(\cdot)$ taking values in the space $X = \{x \in \mathbb{R}_+^R\colon\ \sum_r A_{jr} x_r \le C_j \text{ for all } j \in \mathcal{J}\}$, where $R = |\mathcal{R}|$. (See, for example, Kelly, 1991.)

To make this idea precise, for each $x \in X$, let $m_x(\cdot)$ be the Markov process on E with transition rates given by

$$m \to \begin{cases} m - A_r & \text{at rate } \kappa_r I_{\{m\in\mathcal{A}_r\}} \\ m + A_r & \text{at rate } \mu_r x_r, \end{cases} \tag{1.2}$$

where A_r denotes the vector $(A_{jr}, j \in \mathcal{J})$ and $\infty \pm a = \infty$ for any $a \in \mathbb{Z}_+$. Note that the process $m_x(\cdot)$ is reducible, and so does not always have a unique invariant distribution. Hunt and Kurtz (1994, Theorem 3) show that, provided the distribution of $x^N(0)$ converges weakly to that of $x(0)$, the sequence of processes $x^N(\cdot)$ is relatively compact in $D_{\mathbb{R}^R}[0,\infty)$ and any weakly convergent subsequence has a limit $x(\cdot)$ which obeys the relation

$$x_r(t) = x_r(0) + \int_0^t (\kappa_r \pi_u(\mathcal{A}_r) - \mu_r x_r(u)) du, \tag{1.3}$$

where, for each t, π_t is some invariant distribution of the Markov process $m_{x(t)}(\cdot)$ and additionally satisfies, for all j,

$$\pi_t\{m\colon\ m_j = \infty\} = 1 \text{ if } \sum_{r\in\mathcal{R}} A_{jr} x_r(t) < C_j. \tag{1.4}$$

Thus, at each time t, the invariant distribution π_t acts as a control for the asymptotic process $x(\cdot)$, corresponding to a limiting acceptance rate for calls of each type. For a discussion of this result, which involves a separation, in the limit, of the time scales of the processes $x^N(\cdot)$ and $m^N(\cdot)$, see Hunt and Kurtz (1994) and Bean *et al.* (1995).

Of particular interest is the case where there exists a function π' on X (each value of which is a probability distribution on E) with the property that, for *all* convergent subsequences, we may take $\pi_t = \pi'_{x(t)}$ in eqn (1.3). We may then define a *velocity field* $v = (v_r,\ r \in \mathcal{R})$ on X by

$$v_r(x) = \kappa_r \pi'_x(\mathcal{A}_r) - \mu_r x_r, \tag{1.5}$$

so that eqn (1.3) becomes

$$x_r(t) = x_r(0) + \int_0^t v_r(x(u))du. \tag{1.6}$$

It will then generally be the case that, for all t, $x_r(t)$ is uniquely determined by $x_r(0)$, so that the convergence asserted above takes place in the entire sequence of networks.

Further, when such a velocity field may be defined, it is usually possible to show that, for all t, $x(t)$ is a continuous function of $x(0)$. Since X is compact, the argument of Theorem 3.3 of Bean *et al.* (1995) then applies equally to the present, more general, situation to show that there is at least one *fixed point* $\bar{x} \in X$ such that $v(\bar{x}) = 0$; that is, satisfying the fixed point equations

$$\kappa_r \pi'_x(\mathcal{A}_r) = \mu_r x_r, \quad r \in \mathcal{R}. \tag{1.7}$$

It is scarcely surprising (but for a formal proof see Bean *et al.*, 1994b) that when this fixed point $\bar{x}$ is unique, and further is such that all trajectories of $x(\cdot)$ converge to it, then the invariant distribution of the process $x^N(\cdot)$ converges weakly to the distribution concentrated on the single point $\bar{x}$, while the invariant distribution of the 'free capacity' process $m^N(\cdot)$ converges weakly to $\pi'_{\bar{x}}$. In particular, for each r, $\pi'_{\bar{x}}(\mathcal{A}_r)$ is the limiting equilibrium acceptance probability for calls of type r.

When the process $x(\cdot)$ possesses more than one fixed point, each may be associated, for any large N, with some 'quasi-equilibrium' regime of the process $x^N(\cdot)$, maintained over some extended period of time—as in the example which we discuss in Section 2.

Where there does *not* exist a function π' on X such that $\pi_t = \pi'_{x(t)}$, so that it is impossible to define a velocity field on X, then behaviour in the associated sequence of networks is typically highly pathological. Examples of such behaviour are given by Hunt (1995).

The remainder of this paper is primarily concerned with the identification of conditions under which a velocity field *may* be defined, and with the determination of the resulting dynamics and fixed points of the process $x(\cdot)$. In Section 2 we review results for single resource networks (where a velocity field may always be defined), and in Section 3 we study two-resource networks. Finally, in Section 4 we discuss briefly the general case.

However, it is convenient to make a number of further definitions at this point.

Partition the set X by defining, for each $\mathcal{S} \subseteq \mathcal{J}$, $X_{\mathcal{S}} = \{x \in X\colon \sum_r A_{jr}x_r(t) = C_j$ if and only if $j \in \mathcal{S}\}$. We shall find it convenient to write X_j for $X_{\{j\}}$, and shall make similar obvious notational simplifications elsewhere.

For each subset $\mathcal{S}$ of $\mathcal{J}$, let $E_{\mathcal{S}} = \{m \in E\colon\ m_j < \infty$ if and only if $j \in \mathcal{S}\}$. We assume that the matrix of capacity requirements (A_{jr}) and the acceptance regions $\mathcal{A}_r$ are such that, for each $x \in X$ and $\mathcal{S} \subseteq \mathcal{J}$, there is *at most* a single invariant distribution $\pi_x^{\mathcal{S}}$ of the Markov process $m_x(\cdot)$ on E which assigns probability one to the set $E_{\mathcal{S}}$. (The distribution $\pi_x^{\mathcal{S}}$ may also be thought of as the invariant distribution of the obvious projection of the process $m_x(\cdot)$ onto $\mathbb{Z}_+^{\mathcal{S}}$.) There is no loss of generality in this irreducibility assumption—for a discussion see again Hunt and Kurtz (1994). Note that the distribution $\pi_x^{\emptyset}$ exists for all $x \in X$, assigning probability one to the single point $(\infty, \dots, \infty)$ of the set $E_{\emptyset}$.

Then, from the above results of Hunt and Kurtz, it follows that there exist nonnegative functions $\lambda^{\mathcal{S}}(\cdot)$, $\mathcal{S} \subseteq \mathcal{J}$, summing to one, such that, for almost all t,

$$\pi_t = \sum_{\mathcal{S} \subseteq \mathcal{J}} \lambda^{\mathcal{S}}(t)\pi_{x(t)}^{\mathcal{S}} \tag{1.8}$$

where, from (1.4),

$$\lambda^{\mathcal{S}}(t) = 0 \text{ if } \sum_{r \in \mathcal{R}} A_{jr}x_r(t) < C_j \text{ for any } j \in \mathcal{S}, \tag{1.9}$$

and where additionally we make the convention that $\lambda^{\mathcal{S}}(t) = \lambda^{\mathcal{S}}(t)\pi_{x(t)}^{\mathcal{S}} = 0$ if $\pi_{x(t)}^{\mathcal{S}}$ does not exist. Identification of π_t, $t \geq 0$, thus reduces to identification of the functions $\lambda^{\mathcal{S}}(\cdot)$.

Finally, define also, for each x, each $\mathcal{S} \subseteq \mathcal{J}$ such that $\pi_x^{\mathcal{S}}$ exists, and each $j \in \mathcal{J}$,

$$\alpha_j^{\mathcal{S}}(x) = \sum_{r \in \mathcal{R}} A_{jr}\{\kappa_r\pi_x^{\mathcal{S}}(\mathcal{A}_r) - \mu_r x_r\}. \tag{1.10}$$

The quantity $\alpha_j^{\mathcal{S}}(x)$ will play an important role in subsequent analysis. Note in particular that

$$\alpha_j^{\mathcal{S}}(x) = 0 \text{ if } j \in \mathcal{S}. \tag{1.11}$$

This follows from the observation that, in equilibrium, the jth component of the restriction of the process $m_x(\cdot)$ to $E_{\mathcal{S}}$ has zero drift for each $j \in \mathcal{S}$. A formal proof may be given analogously to that of Lemma 4 of Hunt and Kurtz (1994).

2 Single resource networks

We now consider further the single resource case $\mathcal{J} = \{1\}$. It is convenient to write C for C_1, A_r for A_{1r}, and $\alpha^{\mathcal{S}}(x)$ for $\alpha_1^{\mathcal{S}}(x)$.

Here the compactified space $E = \mathbb{Z}_+ \cup \{\infty\}$ and the requirement that, for each r, the indicator function $I_{\mathcal{A}_r}$ of the acceptance region $\mathcal{A}_r$ be continuous at ∞ implies that there is some *finite* $M \in E$ such that, again for each r, either $m \in \mathcal{A}_r$ for all $m > M$ (including $m = \infty$) or $m \notin \mathcal{A}_r$ for all $m > M$. Define

$\mathcal{R}^* = \{r \in \mathcal{R}\colon\ \infty \in \mathcal{A}_r\}$. Thus $\mathcal{R}^*$ is the set of call types which are accepted for all sufficiently large values of the free capacity in the network.

In most applications we might expect $\mathcal{R}^* = \mathcal{R}$. However, there are practical circumstances where this might not be the case—for example, a call type which was to be allocated less resource when the network was nearly full might be modelled as two call types with disjoint acceptance regions.

Now note that, for all x,

$$\pi_x^{\emptyset}(\mathcal{A}_r) = I_{\{r \in \mathcal{R}^*\}} \tag{2.1}$$

(where again I is the indicator function) and so, from eqn (1.10), $\alpha^{\emptyset}(x) = \sum_r A_r\{\kappa_r I_{\{r\in\mathcal{R}^*\}} - \mu_r x_r\}$. This quantity is also the drift rate *towards* the origin of the process $m_x(\cdot)$ while in the set $[M+1,\infty)$ and so elementary Lyapounov techniques for such processes (see, for example, Fayolle *et al.*, 1995) show that the restriction of this process to $\mathbb{Z}_+$ is ergodic—and so the distribution π_x^1 $(= \pi_x^{\{1\}})$ exists—if and only if $\alpha^{\emptyset}(x) > 0$.

Let $X_1^+ = \{x \in X_1\colon\ \alpha^{\emptyset}(x) > 0\}$ and let $X_1^- = X_1 \setminus X_1^+$. It is then straightforward to show that a velocity field for the limit process $x(\cdot)$ may be defined everywhere on X, the function π'_x being given by

$$\pi'_x = \begin{cases} \pi_x^{\emptyset} & \text{if } x \in X_{\emptyset} \cup X_1^-, \\ \pi_x^1 & \text{if } x \in X_1^+. \end{cases} \tag{2.2}$$

In the case $x \in X_{\emptyset}$ this result follows from eqn (1.4) (or equivalently from eqn (1.9)), while in the case $x \in X_1^-$ it is immediate from the above criterion for the existence of π_x^1. To prove the remaining case note, from eqns (1.3), (1.8), (1.10), and (1.11), together with the condition $\lambda^{\emptyset}(t) + \lambda^1(t) = 1$, we have easily that

$$\sum_{r\in\mathcal{R}} A_r x_r(t) = \sum_{r\in\mathcal{R}} A_r x_r(0) + \int_0^t \lambda^{\emptyset}(u)\alpha^{\emptyset}(x(u))du. \tag{2.3}$$

Since necessarily $\sum_r A_r x_r(t) \le C$ for all t, it follows that $\lambda^{\emptyset}(t) = 0$ for (almost) all t with $x(t) \in X_1^+$.

The simple idea underlying this argument—that the process $x(\cdot)$ must remain within X—is due to Hunt (1990). Hunt and Kurtz (1994) prove the above result in the case $\mathcal{R}^* = \mathcal{R}$. A slightly more formal version of the present argument is given by Bean *et al.* (1994b).

It is now readily verified that, for each r, $\pi'_x(\mathcal{A}_r)$ is Lipschitz continuous on $X_{\emptyset} \cup X_1^-$ (trivially) and also on the set X_1 (see Bean *et al.*, 1995). Hence trajectories of the process $x(\cdot)$ are well-defined functions of their positions at time 0 and discontinuities in the velocity of any trajectory occur only at times of passage from $X_{\emptyset}$ to X_1^+. (Passage from X_1^+ to $X_{\emptyset}$ is impossible by the continuity of the function $\alpha^{\emptyset}$ on X and the relation (2.3)). It follows from standard arguments for dynamical systems that, for each t, $x(t)$ is a continuous function of $x(0)$ and so, as indicated earlier, the process $x(\cdot)$ has at least one fixed point.

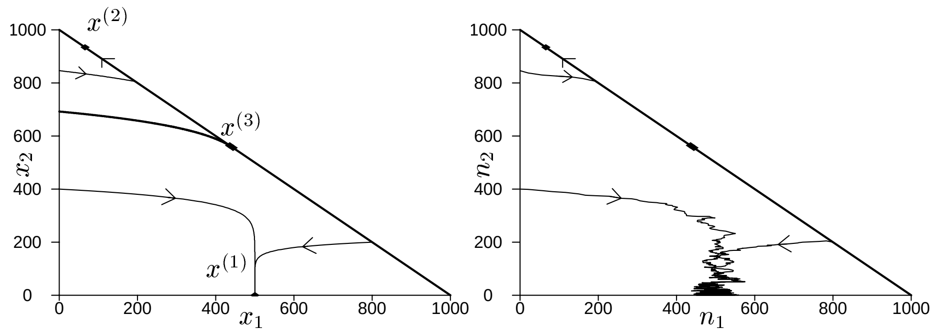

FIG. 1. Analytical and simulation results for numerical example

In the case where

$$\sum_{r\in\mathcal{R}} A_r \kappa_r I_{\{r\in\mathcal{R}^*\}}/\mu_r \le C, \tag{2.4}$$

define $\hat{x} \in X$ by $\hat{x}_r = \kappa_r I_{\{r\in\mathcal{R}^*\}}/\mu_r$. Then $\alpha^{\emptyset}(\hat{x}) = 0$ and so $\hat{x} \in X_{\emptyset} \cup X_1^-$. It follows from eqns (1.7), (2.1), and (2.2) that $\hat{x}$ is the unique fixed point of the process $x(\cdot)$ in the set $X_{\emptyset} \cup X_1^-$. If $\mathcal{R}^* = \mathcal{R}$ then it follows easily from the eqns (1.7) that there is no further fixed point in X^+.

When the condition (2.4) holds but $\mathcal{R}^* \ne \mathcal{R}$, then there may be more than one fixed point. Bean *et al.* (1994b) give a numerical example with two call types. All calls require a single unit of resource and $C = 1000$, $\kappa_1 = 500$, $\kappa_2 = 700$, $\mu_1 = 1.0$, $\mu_2 = 0.1$. The acceptance regions are given by $\mathcal{A}_1 = \{m\colon m > 0\}$ and $\mathcal{A}_2 = \{m\colon 0 < m < 5\}$ (so that here $\mathcal{R}^* = \{1\}$). They show that the process $x(\cdot)$ possesses three distinct fixed points $x^{(1)}$, $x^{(2)}$, $x^{(3)}$. Every trajectory of the process $x(\cdot)$ tends to one of these points, although the point $x^{(3)}$ is unstable in the sense of possessing a domain of attraction of Lebesgue measure zero in X. The limit behaviour of the corresponding sequence of networks is therefore essentially bistable. The left panel of Fig. 1 shows sample trajectories of the process $x(\cdot)$—the thick line separates the domains of attraction of $x^{(1)}$ and $x^{(2)}$ and is of course itself a trajectory of the system, tending to $x^{(3)}$. The right panel shows simulated trajectories of the process $x^1(\cdot)$ $(= n^1(\cdot))$ in the associated sequence of networks. Here C is sufficiently large that the process $x^1(\cdot)$ should be reasonably well-approximated by $x(\cdot)$ and indeed the bistable behaviour of $x^1(\cdot)$ is clearly evident. However, this process is of course ergodic, so that, over sufficiently long time periods, it alternates between typically lengthy residences in the neighbourhoods of $x^{(1)}$ and $x^{(2)}$.

In the case where the relation (2.4) does *not* hold the fixed points of the process $x(\cdot)$ necessarily lie in X_1^+. Where, additionally, $A_r = 1$ for all r, an argument of Bean *et al.* (1995) for the case $\mathcal{R}^* = \mathcal{R}$ extends unchanged to the present case to show that there is a unique fixed point $\bar{x} \in X_1^+$. Provided only that all trajectories of $x(\cdot)$ then converge to $\bar{x}$ (this is difficult to show formally except in the case $R = 2$), identification of this point via the equations (1.7) permits the determination of limiting equilibrium behaviour—in particular limiting call acceptance probabilities—for the associated sequence of networks.

3 Two-resource networks

We now study the two-resource case $\mathcal{J} = \{1, 2\}$. Here some distinctly pathological behaviour is possible, as is shown by the example of Hunt (1995), which we discuss briefly below. We require conditions under which such pathological behaviour may not occur.

For any x and j, the restriction of the process $m_x(\cdot)$ to E_j is essentially one-dimensional and it follows, as in the previous section, that the distribution π_x^j exists if and only if $\alpha_j^{\emptyset}(x) > 0$. It again follows as there, and using the condition (1.4), that when $x(t) \notin X_{12}$ then $\pi_t = \pi'_{x(t)}$ where π'_x is given by

$$\pi'_x = \begin{cases} \pi_x^{\emptyset} & \text{if } x \in X_{\emptyset} \cup X_1^- \cup X_2^-, \\ \pi_x^j & \text{if } x \in X_j^+, \end{cases} \tag{3.1}$$

and where, for each j, $X_j^+ = \{x \in X_j \colon \alpha_j^{\emptyset}(x) > 0\}$, $X_j^- = X_j \setminus X_j^+$. It remains to consider the identification of π_t in the case where $x(t) \in X_{12}$. The key here is again given by the functions $\alpha_j^{\mathcal{S}}$.

For either $j \in \mathcal{J}$, let j' denote its complement in $\mathcal{J}$. For each j, define the function β_j on X by

$$\beta_j(x) = \begin{cases} \alpha_j^{j'}(x) & \text{if } \alpha_{j'}^{\emptyset}(x) > 0, \\ \alpha_j^{\emptyset}(x) & \text{if } \alpha_{j'}^{\emptyset}(x) \le 0. \end{cases} \tag{3.2}$$

Recall that $\alpha_j^{j'}(x)$ is defined if and only if $\alpha_{j'}^{\emptyset}(x) > 0$. The quantity $\beta_j(x)$ also has an informal interpretation in terms of the restriction of the process $m_x(\cdot)$ to $E_{12} = \mathbb{Z}_+^2$. In the case $\alpha_{j'}^{\emptyset}(x) > 0$, suppose that the component j of this restricted process is far from 0 but the component j' is in equilibrium; then $\beta_j(x)$ is the averaged (negative) drift rate of the component j. In the case $\alpha_{j'}^{\emptyset}(x) \le 0$, a similar but simpler interpretation holds. These ideas may be formalized as, for example, by Fayolle *et al.* (1995), but for our purposes a formal definition is more easily made as above in terms of the invariant distributions associated with the restrictions of the process $m_x(\cdot)$ to $E_{j'}$ or $E_{\emptyset}$ as appropriate.

Define subsets of X_{12} as follows. Let

$$\begin{aligned} U &= \{x \in X_{12} \colon \beta_1(x) \wedge \beta_2(x) > 0\}, \\ V_j &= \{x \in X_{12} \colon \beta_j(x) > 0,\ \beta_{j'}(x) \le 0\}, \quad j = 1, 2, \\ W &= \{x \in X_{12} \colon \beta_1(x) \vee \beta_2(x) \le 0\}, \\ W^- &= \{x \in W \colon \alpha_1^{\emptyset}(x) \vee \alpha_2^{\emptyset}(x) \le 0\}, \end{aligned}$$

$$W^+ \quad = \quad \{x \in W\colon\ \alpha_1^{\emptyset}(x) \wedge \alpha_2^{\emptyset}(x) > 0\}.$$

Note that it follows from the definition (3.2) that $W = W^- \cup W^+$, so that the above sets form a partition of X_{12}.

Bean *et al.* (1994b) show that, under the condition

$$(\infty, \infty) \in \mathcal{A}_r \text{ for all } r \in \mathcal{R}, \tag{3.3}$$

for (almost) all t,

$$\pi_t = \begin{cases} \pi^{12}_{x(t)} & \text{if } x(t) \in U, \\ \pi^{j}_{x(t)} & \text{if } x(t) \in V_j, \quad j = 1, 2, \\ \pi^{\emptyset}_{x(t)} & \text{if } x(t) \in W^-. \end{cases} \tag{3.4}$$

We shall discuss below the necessity of the condition (3.3) for the result (3.4), and also the remaining case, $x(t) \in W^+$. However, note that when the result (3.4) holds and W^+ is empty (as will usually be the case in applications), it is again possible to define a velocity field for the process $x(\cdot)$ everywhere on X.

The result (3.4) is proved using essentially the same arguments as those used to establish the result (2.2) in the single resource case. We give here an outline. Note first that, under the condition (3.3), it follows from the definitions (1.10) and (3.2), that, for all x,

$$\beta_j(x) \le \alpha_j^{\emptyset}(x), \quad j = 1, 2. \tag{3.5}$$

Further, again under this condition (3.3), standard results for Markov chains on $\mathbb{Z}_+^2$ with partial spatial homogeneity (see Fayolle *et al.*, 1995, or Zachary, 1995) show that, for all x,

$$\pi_x^{12} \text{ exists if and only if } \beta_1(x) \wedge \beta_2(x) > 0. \tag{3.6}$$

Note also that, analogously to eqn (2.3), and by again using in particular the result (1.11), we have that for each j,

$$\sum_{r \in \mathcal{R}} A_{jr} x_r(t) = \sum_{r \in \mathcal{R}} A_{jr} x_r(0) + \int_0^t \{\lambda^{\emptyset}(u) \alpha_j^{\emptyset}(x(u)) + \lambda^{j'}(u) \alpha_j^{j'}(x(u))\} du. \tag{3.7}$$

From eqn (3.5), for t with $x(t) \in U$ and each j, $\alpha_j^{\emptyset}(x(t)) > 0$ and $\alpha_j^{j'}(x(t)) > 0$. It follows from eqn (3.7), arguing as in the single resource case, that, for (almost) all t with $x(t) \in U$, $\lambda^{\emptyset}(t) = \lambda^1(t) = \lambda^2(t) = 0$ and so $\pi_t = \pi^{12}_{x(t)}$ as required.

The remaining cases of the result (3.4) are proved similarly, on making use also of the result (3.6).

It seems likely that the result (3.6) continues to hold in the absence of the condition (3.3) (here the only doubt in the existing literature lies with the boundary case $\beta_1(x) \wedge \beta_2(x) = 0$) in which case a relatively straightforward variation of the above argument may be used to show that the result (3.4) also continues to hold. Thus, under what are at worst mild regularity conditions, and certainly in the case where the condition (3.3) does hold, a velocity field for the process $x(\cdot)$

may be defined everywhere on the set $X \setminus W^+$. When W^+ is empty, again only mild regularity conditions are required to show that the trajectories of the process $x(\cdot)$ are well-defined and that, for each t, $x(t)$ is a continuous function of $x(0)$. It follows that, in this case, there exists at least one fixed point for the process $x(\cdot)$.

Hunt (1995) gives an example in which the set W^+ is nonempty. Here, and in general, for t such that $x(t) \in W^+$,

$$\pi_t = \lambda^1(t)\pi^1_{x(t)} + \lambda^2(t)\pi^2_{x(t)}, \tag{3.8}$$

where as usual $\lambda^1(t)$ and $\lambda^2(t)$ are positive and sum to one. However, beyond this, the behaviour of the process $x(\cdot)$ within the set W^+ is indeterminate, corresponding to the fact that here the sequence of processes $x^N(\cdot)$ may have different limits in different subsequences. Trajectories of two such limits may agree up to the time of entrance into W^+, but behave quite differently thereafter.

It is therefore important for the control of networks to have conditions which ensure that the set W^+ is empty. For each $r \in \mathcal{R}$, let $\mathcal{J}_r = \{j \in \mathcal{J}\colon\ A_{jr} > 0\}$. Extend the definition of $\mathcal{R}^*$ given in the previous section to two- (and more) resource networks by letting

$$\mathcal{R}^* = \{r \in \mathcal{R}\colon\ E_{\mathcal{S}} \subset \mathcal{A}_r \text{ for all } \mathcal{S} \text{ with } \mathcal{S} \cap \mathcal{J}_r = \emptyset\}. \tag{3.9}$$

Thus, using also the continuity of $I_{\mathcal{A}_r}$, calls of type $r \in \mathcal{R}^*$ are accepted for all sufficiently large values of the free capacities of those resources in $\mathcal{J}_r$—regardless of the state of the remaining resources. In a variation of Conjecture 5 of Hunt and Kurtz (1994) we conjecture that a sufficient condition for W^+ to be empty is given by $\mathcal{R}^* = \mathcal{R}$. (This of course implies in particular the condition (3.3).)

The following theorem shows this to be the case where $A_{1r} = A_{2r}$ for those call types r such that $A_{1r} \wedge A_{2r} > 0$. It generalizes a result of Moretta (1995).

Theorem 3.1 *Suppose that $\mathcal{R}^* = \mathcal{R}$ and that*

$$A_{1r} = A_{2r} \textit{ for all } r \textit{ with } \mathcal{J}_r = \mathcal{J}. \tag{3.10}$$

Then W^+ is empty.

Proof Suppose there exists $x \in W^+$. Then, for each j, $\alpha^{\emptyset}_j(x) > 0$ and so the distribution π^j_x exists. Thus, again for each j, $\alpha^{j'}_j(x) = \beta_j(x) \le 0$, and since also (by the result (1.11)) $\alpha^j_j(x) = 0$, it follows that

$$\sum_{r\in\mathcal{R}} A_{jr}\kappa_r\{\pi^{j'}_x(\mathcal{A}_r) - \pi^j_x(\mathcal{A}_r)\} \le 0. \tag{3.11}$$

For r such that $A_{jr} > 0$, either $\mathcal{J}_r = \{j\}$ or $\mathcal{J}_r = \mathcal{J}$. In the former case the condition $\mathcal{R}^* = \mathcal{R}$ implies that $\pi^{j'}_x(\mathcal{A}_r) = 1$, while $\pi^j_x(\mathcal{A}_r) < 1$ (since necessarily $E_j \not\subseteq \mathcal{A}_r$). Hence

$$\sum_{r\colon\ \mathcal{J}_r=\mathcal{J}} A_{jr}\kappa_r\{\pi^{j'}_x(\mathcal{A}_r) - \pi^j_x(\mathcal{A}_r)\} \le 0 \tag{3.12}$$

with strict inequality if $\mathcal{J}_r = \{j\}$ for at least one r.

The irreducibility assumption of Section 1 implies that there is at least one call type r such that $\mathcal{J}_r = \{1\}$ or $\mathcal{J}_r = \{2\}$. It follows, on interchanging j and j' in eqn (3.12), and using the condition (3.10), that the eqn (3.12) is self-contradictory. □

As usual the fixed points of the process $x(\cdot)$ are determined by solution of the equations (1.7). In the case where, for either j, $\sum_r A_{jr}\kappa_r/\mu_r \leq C_j$, we may in effect replace C_j by ∞ and consider the single resource j'. Otherwise, and when the set U is nonempty, the analysis may be more complicated, involving in particular the (nontrivial) determination of two-dimensional invariant distribution π_x^{12} for $x \in U$.

We know of no example in which $\mathcal{R}^* = \mathcal{R}$ and there is more than one fixed point. Moretta (1995) considers the case where $\mathcal{R}^* = \mathcal{R}$ and the matrix $(A_{jr},\ j \in \mathcal{J},\ r \in \mathcal{R})$ is given by

$$A_{jr} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

He uses a coupling argument to show that here, if there is more than one fixed point, then all fixed points necessarily lie in U. If therefore an (essentially straightforward single resource) analysis identifies a fixed point outside this set, there will be no further fixed point within it. Moretta also presents compelling evidence that, for this model, there is only ever one fixed point.

Moretta also considers the problem of the determination of the invariant distribution π_x^{12} for $x \in U$, and that of determining more refined approximations to call acceptance probabilities in networks whose capacities are insufficiently large to justify direct application of the above asymptotic theory.

4 General networks

In the previous two sections we have outlined an essentially complete theory for the identification of the 'driving' distribution π_t of eqn (1.3) in the case of single and two-resource networks. This has used little more than Hunt's elementary observation that the process $x(\cdot)$ must remain within the set X. (Only for t such that $x(t)$ belongs to the set W^+, defined in the previous section, is a more careful argument required, and this too is due to Hunt (1995).)

For networks with more than two resources, the identification of π_t is very much more complex. For $x \in X$, define a set $\mathcal{S} \subseteq \mathcal{J}$ to be *blocking* with respect to x if $\pi_x^{\mathcal{S}}$ exists and $\sum_r A_{jr}x(t)_r = C_j$ for all $j \in \mathcal{S}$. One very reasonable conjecture is that, for any t, $\pi_t = \pi_{x(t)}^{\mathcal{S}}$ whenever there exists a 'maximal' blocking set $\mathcal{S}$ with respect to $x(t)$ containing every other such blocking set.

Again as remarked earlier, we are particularly interested in the identification of conditions under which a velocity field may be defined for the process $x(\cdot)$. We hesitate to make any conjectures here, but merely observe that for none of the 'pathological' examples of Hunt (1995) is the condition $\mathcal{R}^* = \mathcal{R}$ satisfied.

Acknowledgements The author is grateful to many people for their contributions to this paper—especially to Nigel Bean, Richard Gibbens, Phil Hunt, Frank Kelly and Brian Moretta, and to Edward Ionides for assistance in constructing the simulations.

Bibliography

1. Bean, N.G., Gibbens, R.J., and Zachary, S. (1994a). The performance of single resource loss systems in multiservice networks. In Jacques Labetoulle and James W. Roberts (eds), *The Fundamental Role of Teletraffic in the Evolution of Telecommunications Networks*, Proceedings of the 14th International Teletraffic Congress, pp. 13–21, Elsevier Science B.V.
2. Bean, N.G., Gibbens, R.J., and Zachary, S. (1994b). Dynamic and equilibrium behaviour of controlled loss networks. Research Report #94-31, Statistical Laboratory, The University of Cambridge, 16 Mill Lane, Cambridge, CB2 1SB, U.K.
3. Bean, N.G., Gibbens, R.J., and Zachary, S. (1995). Asymptotic analysis of large single resource loss systems under heavy traffic, with applications to integrated networks. *Adv. Appl. Probab.*, **27**, 273–292.
4. Fayolle, G., Malyshev, V.A., and Menshikov, M.V. (1995). *Topics in the Constructive Theory of Countable Markov Chains.* Cambridge University Press.
5. Hunt, P.J. (1990). *Limit theorems for stochastic loss networks.* Ph.D. dissertation, University of Cambridge.
6. Hunt, P.J. (1995). Pathological behaviour in loss networks. *J. Appl. Probab.*, **32**, 519–533.
7. Hunt, P.J. and Kurtz, T.G. (1994). Large loss networks *Stochastic Process. Appl.*, **53**, 363–378.
8. Kelly, F.P. (1991). Loss networks. *Ann. Appl. Probab.* **1**, 319–378.
9. Moretta, B. (1995). *Behaviour and control of single and two resource loss networks.* Ph.D. dissertation, Heriot-Watt University.
10. Zachary, S. (1995). On two-dimensional Markov chains in the positive quadrant with partial spatial homogeneity. *Markov Process. Relat. Fields*, **1**, 267–280.