THRESHOLD PHENOMENA IN THE TRANSIENT BEHAVIOUR OF MARKOVIAN MODELS OF COMMUNICATION NETWORKS AND DATABASES

V. ANANTHARAM *

School of Electrical Engineering, Cornell University, Ithaca, NY 14853, U.S.A.

Received 4 August 1988; revised 3 March 1989

Abstract

This paper accompanies a talk given at the Workshop on Mathematical Methods in Queueing Networks held at the Mathematical Sciences Institute at Cornell University in August 1988. In earlier work we had exhibited a threshold phenomenon in the transient behaviour of a closed network of ./M/1 nodes: When there are N customers circulating, and the initial state is x, let $d_x^N(t)$ denote the total variation distance between the distribution at time t and the stationary distribution. Let $d^N(t) = \max_x d_x^N(t)$. We explicitly found a_N proportional to N such that $d^N(ta_N) \to 1$ for *every* $t < 1$, and $d^N(ta_N) \to 0$ for *every* $t > 1$. Thus it appears that the network has not yet converged to stationarity upto a_N, but has converged to stationarity after a_N, so a_N can be naturally interpreted as the settling time of the network. Here we briefly deal with some other similar models – closed networks of ./M/m nodes, a well studied model for circuit switched networks, and a model of Mitra for studying concurrency control in databases. Similar threshold phenomena are established in the transient behaviour of these models.

Keywords: Circuit switching, databases, Jackson networks, Markov chains, queueing, transient behaviour.

1. Introduction. Closed networks of ./M/1 nodes

The introduction of national and global communication networks, the widespread use of local area networks and cellular radio, the need for large databases, the trend towards the development of parallel computers with large numbers of moderately powerful processors, and the importance of improved manufacturing techniques have all lead to a remarkable resurgence of interest in queueing theory. The emphasis in the modern work tends to be on networks, on control issues, approximations and numerical methods, and on asymptotics, e.g. light and heavy traffic limits. Further the systems of practical interest are often

* Research supported by the National Science Foundation, Grant No. NCR 8710840.

large scale systems, in the sense that the number of packets, number of nodes, buffer sizes etc. are large according to common usage. This has lead to a considerable interest in simulation based performance analysis, and design schemes using simulation based sensitivity analysis.

An important issue in simulation is the so called *initial bias problem*, Bratley et al., [6], chapter 3. This is the problem of how long one needs to wait before a simulation run can be considered sufficiently "warmed up", i.e. generating typical statistics. To understand this time in any specific system it is necessary to understand the transient behaviour of the system. It turns out that the systems commonly encountered are such that it is possible to give a very precise estimate for this time. This is the content of some of our earlier work, [4], [5], and the results described in this paper.

Consider a general irreducible continuous time Markov chain $(X(t), t \geqslant 0)$ with finite state space X and stationary distribution π. We are interested in studying the transient behaviour of this chain. It is well known that whatever initial distribution the chain is started in, its distribution will eventually converge to the stationary distribution in most senses of interest. Suppose the chain is started in state $x \in X$. A natural measure of the distance from stationarity at the time t is the total variation distance

$$d_x(t) = \tfrac{1}{2} \sum_{y \in X} | P_x(X(t) = y) - \pi(y) |$$

between the distribution of the state at time t and the stationary distribution. The worst case distance from stationarity over all initial conditions, measured by

$$d(t) = \max_{x \in X} d_x(t)$$

then provides a reasonable description of the transient behaviour of the chain. One knows that $d(t) \to 0$ as $t \to \infty$. In fact, if λ_2 denotes the second eigenvalue of the time 1 transition matrix of the Markov chain, one also knows that

$$d(t) \sim C\lambda_2^t \quad \text{as } t \to \infty,$$

for some constant C. Traditionally one thinks of λ_2 as describing the speed with which the chain approaches stationarity and of $-(\log \lambda_2)^{-1}$ as a time constant describing how long one has to wait before the chain has essentially converged.

The difficulty with considering this description satisfactory is that it is concerned only with the tail behavior of the convergence, and ignores the finite time behavior. This was first convincingly demonstrated by David Aldous, Persi Diaconis and their co-workers in a number of examples, [1], [2], [3], [9]. If λ_2 has large multiplicity, the qualitative behavior of $d(t)$ is somewhat like the graph in fig. 1. In particular, there appears to be a time τ such that $d(t)$ is very close to 1 upto τ, decreases rapidly near τ and then stays close to 0, with the tail being described in the traditional way by the second eigenvalue λ_2. It is natural to think of τ as being the time at which the chain has "become stationary".

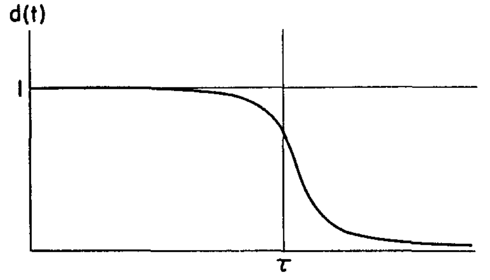

Fig. 1.

This paper is concerned with demonstrating the existence of a phenomenon of this type in the transient behaviour of several commonly studied models of communication networks and databases.

To give a concrete example right away, we will now describe the transient behaviour of closed networks of ./M/1 nodes. Consider a closed network consisting of J exponential servers with service rates $\mu_1, \dots, \mu_J$ and irreducible routing matrix $R = [r_{ij}]$, and with N customers circulating in the network. Recall that this means that a customer in queue at node i waits until all customers ahead of him have been served, and then receives service whose duration is an exponential random variable of mean μ_i^{-1}, all service times being independent. A customer finishing service at node i is routed to node j with probability r_{ij}, all routing decisions being independent and independent of the service times. Further, R is a nonnegative irreducible stochastic matrix. See fig. 2.

The network may be described by a Markov process $\boldsymbol{x} = (X(t),\ t \geqslant 0)$ with state space

$$X = \left\{ \boldsymbol{n} = (n_1, \dots, n_J):\ \sum_{i=1}^{J} n_i = N \right\},$$

where the state $\boldsymbol{n}$ corresponds to there being n_i customers in queue at node i. The

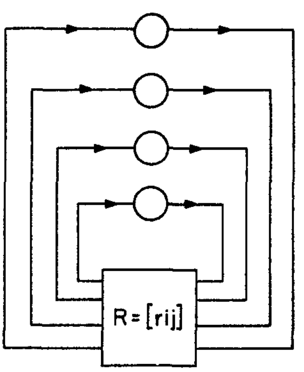

Schematic representation of a network of ·/M/I nodes

Fig. 2.

stationary distribution of the network is well known, see e.g. Gordon and Newell, [10], Kelly, [14], or Walrand, [22], and is given by

$$\pi(\boldsymbol{n}) = c\prod_{i=1}^{J}\left(\frac{\lambda_i}{\mu_i}\right)^{n_i},$$

where c is a normalizing constant and λ_i, $1 \leqslant i \leqslant J$, are positive solutions of the flow balance equations:

$$\lambda_i = \sum_{j=1}^{J} \lambda_j r_{ji}, \quad i = 1,\dots, J.$$

Consider the J numbers μ_j/λ_j, $1 \leqslant j \leqslant J$. When suitably scaled, μ_j/λ_j has the interpretation of the effective service rate at node j. Assume that there is a unique node with smallest effective service rate. Denote this node by s. Let α_{ij} denote the probability that a customer leaving node i will meet node j at least once before returning to node i. Let w denote any node such that

$$\mu_w\alpha_{ws} - \mu_s\alpha_{sw} = \min_{i\neq s}(\mu_i\alpha_{is} - \mu_s\alpha_{si}).$$

Note that $\mu_w\alpha_{ws} - \mu_s\alpha_{sw} > 0$. Indeed, to see this it is enough to observe that $\lambda_i\alpha_{is} = \lambda_s\alpha_{si}$ for all $i \neq s$. Let

$$a_N = \frac{N}{\mu_w\alpha_{ws} - \mu_s\alpha_{sw}}.$$

Let

$$d_x^N(t) = \tfrac{1}{2}\sum_{y\in X} |P_x(X(t) = y) - \pi(y)|,$$

and let $d^N(t) = \max_{x\in X} d_x^N(t)$ measure the worst case distance from the stationary distribution over all possible initial conditions, when there are N customers circulating in the network. Then we have the following

THEOREM 1.1

For every $t < 1$

$$\lim_{N\to\infty} d^N(ta_N) = 1, \tag{1.1}$$

whereas, for every $t > 1$

$$\lim_{N\to\infty} d^N(ta_N) = 0. \tag{1.2}$$

In the sense of theorem 1.1, in the large N limit, the network has not yet converged to stationarity upto a_N but has converged to stationarity after a_N. This suggests that a_N should be considered the correct notion of settling time for the network.

Full details of the proof of this theorem are available in Anantharam, [5], and a more rapid summary of the proof is available in Anantharam, [4]. In section 2 we indicate a general framework along which to build proofs of results such as theorem 1.1. The approach is that proposed by Aldous, [1]. In section 3, we will state a similar theorem for networks of ./M/m nodes and give a rapid sketch of its proof following the lines of our earlier proof of theorem 1.1. This may also help to understand theorem 1.1. In sections 4 and 5 resp. we will show that similar threshold phenomena occur in a well studied Markovian model for circuit switched networks. Brockmeyer et al., [7], Burman et al., [8], Kelly, [12]–[15], Lagarias et al., [16], and Ziednis, [23], and in a model for studying database concurrency control due to Mitra, [17]. The proofs here are relatively straightforward and the basic intuition more transparent. Finally, we will make some concluding remarks in section 6.

Throughout the paper $>_s$ is used to denote stochastic ordering between random variables, and $=^d$ is used to denote equality in distribution.

2. Skeleton proof technique

In this section we outline the skeletal structure of the proofs of the succeeding sections. The methodology is that suggested by Aldous, [1]. We use the notation of theorem 1.1, so the space parameter is N. In section 4, the notation is somewhat different.

To identify a threshold time, we first prove the lower bound (1.1) and then the upper bound (1.2). To prove the lower bound, fix $\delta > 0$, however small. We find a subset $Y \subseteq X$ such that $\pi(Y) \geqslant 1 - \delta$ and an element $x \in X$ such that the time to hit Y starting from x is large. Namely, if $T = \inf\{t > 0\colon X(t) \in Y\}$, we prove that, for every $t < 1$

$$\lim_{N \to \infty} P_x(T \leqslant ta_N) = 0. \tag{2.1}$$

Note that $P_x(T \leqslant ta_N) \geqslant P_x(X(ta_N) \in Y)$. Hence, if we can prove (2.1), it follows from the definitions that $d^N(ta_N) \geqslant d_x^N(ta_N) \geqslant |P_x(X(ta_N) \in Y) - \pi(Y)|$. Hence

$$\lim_{N \to \infty} d^N(ta_N) \geqslant 1 - \delta.$$

Letting $\delta \to 0$ establishes the lower bound.

To prove the upper bound, for any pair of initial states, x, $y \in X$, we couple the paths of the network started from x to those of the network started from y. This means that on the same sample space we construct two X valued Markov processes $(X^A(t),\ t \geqslant 0)$ and $(X^B(t),\ t \geqslant 0)$ such that $X^A(0) = x$ and $X^B(0) = y$ and such that (1) the A (respectively B) process is the same in law as the state process of the network started at x (respectively y) and (2) there is a proper random time T_{xy} after which the A and B processes agree, i.e.

$$X_i^A(t) = X_i^B(t) \quad \text{for all } 1 \leqslant i \leqslant J \quad \text{and } t \geqslant T_{xy}.$$

Note that we do not require that T_{xy} is the first time at which the A and B processes agree. We call T_{xy} the coupling time. For more on coupling, see e.g. Griffeath, [11], and Ross, [20].

Now let

$$\rho_{xy}^N = \tfrac{1}{2} \sum_{z \in X} | P_x(X(t) = z) - P_y(X(t) = z) |.$$

Then it is easy to see that $d^N(t) \leqslant \sup_{xy} \rho_{xy}^N(t)$. Further, for a coupling time T_{xy} it is easy to see that, for any $t \geqslant 0$, $\rho_{xy}^N(t) \leqslant P(T_{xy} \leqslant t)$. Now suppose that we can construct couplings for each pair of states x, $y \in X$ such that, for every $t > 1$

$$\lim_{N \to \infty} \sup_{xy} P(T_{xy} \leqslant t a_N) = 0.$$

Then by the above, we have, for every $t > 1$,

$$\lim_{N \to \infty} \sup_{xy} \rho_{xy}^N(t a_N) = 0.$$

Hence, for every $t > 1$,

$$\lim_{N \to \infty} d^N(t a_N) = 0.$$

This establishes the upper bound.

3. Closed networks of ./M/m nodes

In this section we will sketch a proof of the generalization of Thm. 1.1 to a closed network of /M/m nodes, emphasizing only how to modify the proof of Thm. 1.1 which was given in [5].

Consider a closed network of J nodes with node j being an m_j server node with exponential servers having service rate μ_j. Let the routing matrix be denoted $R = [r_{ij}]$, and let there be N customers circulating in the network. See fig. 3. The network may be described by a Markov process $\boldsymbol{x} = (X(t),\ t \geqslant 0)$ with state space

$$X = \left\{ \boldsymbol{n} = (n_1, \ldots, n_J) : \sum_{i=1}^{J} n_i = N \right\},$$

where the state $\boldsymbol{n}$ corresponds to there being n_i customers in queue at node i. The stationary distribution of the network is well known, see e.g. Gordon and Newell, [10], Kelly, [14], or Walrand, [22], and is given by

$$\pi(\boldsymbol{n}) = c \prod_{i=1}^{J} \alpha_i(n_i) \left(\frac{\lambda_i}{m_i \mu_i} \right)^{n_i},$$

where c is a normalizing constant, and λ_i, $1 \leqslant i \leqslant J$, are positive solutions of the flow balance equations:

$$\lambda_i = \sum_{j=1}^{J} \lambda_j r_{ji}, \quad i = 1, \ldots, J.$$

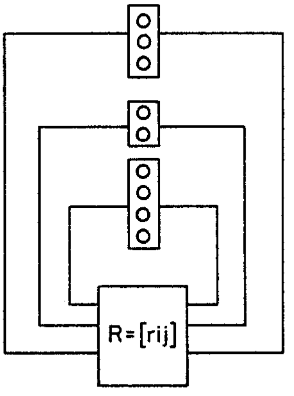

Schematic representation of a network of •/M/m nodes

Fig. 3.

Here

$$\alpha_i(n_i) = \frac{m_i^{n_i}}{n_i!} \quad \text{if } n_i < m_i,$$

and

$$\alpha_i(n_i) = \frac{m_i^{m_i}}{m_i!} \quad \text{if } n_i \geqslant m_i.$$

When suitably scaled, $m_j\mu_j/\lambda_j$ has the interpretation of the effective service rate at node j. Assume that there is a unique node with smallest effective service rate and denote this node by s. Let α_{ij} denote the probability that a customer leaving node i will meet node j at least once before returning to node i. Let w denote any node such that

$$m_w\mu_w\alpha_{ws} - m_s\mu_s\alpha_{sw} = \min_{i \neq s}(m_i\mu_i\alpha_{is} - m_s\mu_s\alpha_{si}).$$

Let

$$\alpha_N = \frac{N}{m_w\mu_w\alpha_{ws} - m_s\mu_s\alpha_{sw}}.$$

Let

$$d_x^N(t) = \tfrac{1}{2}\sum_{y \in X} |P_x(X(t)=y) - \pi(y)|$$

and let $d^N(t) = \max_{x \in X} d_x^N(t)$ measure the worst case distance from the stationary distribution over all possible initial conditions when there are N customers circulating in the network. Then we have the following

THEOREM 3.1

For every $t < 1$

$$\lim_{N \to \infty} d^N(ta_N) = 1, \tag{3.1}$$

whereas, for every $t > 1$

$$\lim_{N \to \infty} d^N(ta_N) = 0. \tag{3.2}$$

Proof

It is easy to see from the form of the stationary distribution that for any $\delta > 0$ we can find $K > 0$ independent of N such that, if $Y = \{\boldsymbol{n} \in X : n_s > N - K\}$ then $\pi(Y) > 1 - \delta$ for all sufficiently large N. From the general reasoning of section 2, in order to prove the lower bound, (3.1), it suffices to find an initial condition x such that if T denotes the time to hit Y starting from x, then $\lim_{N \to \infty} P_x(T > ta_N) = 1$ for every $t < 1$. As in the proof of Thm. 1.1. in Anantharam, [5], we take x to be the initial condition in which $N^{4/5}$ of the packets are at node s and the remaining $N - N^{4/5}$ are at node w. We set the closed network $\boldsymbol{N}$ up on a sample space supporting independent Poisson processes $V_j = (V_j(t),\ t \geqslant 0)$ of rates $m_j \mu_j$, $1 \leqslant j \leqslant J$, respectively, (the virtual departure processes), i.i.d. families $C_j = (C_j(n),\ n \geqslant 0)$, $1 \leqslant j \leqslant J$, with each $C_j(n)$ being uniformly distributed on $\{1, \ldots, m_j\}$ and also routing prescriptions for each virtual departure at the individual nodes, which are independent, independent of the V_j and C_j, $1 \leqslant j \leqslant J$, and appropriately distributed. The paths are constructed as follows: when the nth point of V_j occurs, look at $C_j(n)$. Let x_j denote the number of packets in node j at this time. If $C_j(n) \leqslant x_j$, release any one of the packets in node j, and decide where it goes based on what the nth routing prescription at node j says. If $C_j(n) > x_j$ do nothing. An important point to note is that the process of queue size at the nodes is not affected by which packet moves. This allows us to prove several intuitively obvious statements rigorously using elementary couplings. We learnt this idea from Van der Wal, [21].

For example, let $\boldsymbol{N}_w$ denote the portion of the original network other than node w thought of as an open network and fed by the virtual departure process of node w. Start $\boldsymbol{N}_w$ with $N^{4/5}$ customers at node s and the other nodes empty. We may set $\boldsymbol{N}_w$ up on the same sample space that we used for $\boldsymbol{N}$ in the preceding paragraph. Let T denote the first time at which the total number of customers in $\boldsymbol{N}_w$ becomes $N - K$. Then we claim $T_1 < T$. The idea is the following–suppose V_w has its nth point, let x_w denote the number of customers at node w in $\boldsymbol{N}$ at this time, and suppose $C_w(n) > x_w$. Then of course there is no movement in $\boldsymbol{N}$, however there is an arrival to $\boldsymbol{N}_w$. We color this arrival red. We can arrange internal movements so that the process of uncolored packet sizes at the nodes evolves like the queue sizes process in $\boldsymbol{N}$, while the process of total number of uncolored and colored packets at the nodes evolves like the queue sizes process in

N_w. Clearly, this would prove $T_1 < T$. To arrange the desired situation, at the nth point of V_j, $j \neq w$, we look at $C_j(n)$. Let x_j denote the number of uncolored packets and x_j^w the number of colored and uncolored packets in node j at this time. If $C_j(n) \leqslant x_j$ and $C_j(n) \leqslant x_j^w$, we move an uncolored packet out of j. (Since $C_j(n) \leqslant x_j$ implies $x_j \geqslant 1$, there must be at least one such packet.) If $C_j(n) > x_j$ but $C_j(n) \leqslant x_j^w$, we move a red packet out of node j. (Since $x_j^w \geqslant C_j(n) > x_j$, there must be at least one such packet.) If $C_j(n) > x_j$ and $C_j(n) > x_j^w$, we do nothing. This gives the desired. A formal proof is by induction on the virtual service times.

A comparison of the above with the corresponding part of [5] shows that we have basically had to be just a little bit more careful in the coupling.

Let N_{sw} denote the portion of the network other than the nodes s and w, started initially empty and fed by the virtual departure processes of these nodes. From lemma 5.1 of [5] it follows easily that N_{sw} is a *critical* open network, i.e. that the input rate to any of nodes $j \neq s$, w is at most $m_j \mu_j$. One can then prove, exactly like lemma A.3 of [5], that on a time scale of order $N^{4/3}$ the total number of customers in N_{sw} does not exceed $N^{3/4}$ with probability tending to 1. Namely, if $(X_j(t),\ t \geqslant 0)$ denotes the queue size process at node $j \neq s$, w in N_{sw}, then

$$\lim_{N \to \infty} P\left(\sum_{j \neq s, w} X_j(t) \leqslant N^{3/4} \text{ for all } t \leqslant N^{4/3} \right) = 1. \tag{3.3}$$

Now let us realize N_w on a sample space supporting V_j and C_j, $1 \leqslant j \leqslant J$ and two families of i.i.d. routing prescriptions–for each virtual departure from node w, its entire path through N_w till it first visits s or leaves the network is decided, and similarly for each departure from node s. This sample space can support precisely those initial conditions for N_w where all the initial customers are at node s. Let $(\hat{X}_j(t),\ t \geqslant 0)$ denote the queue size process at node $j \neq w$ when N_w is started with $N^{4/5}$ customers at node s and with FCFS departures at the individual nodes. Let

$$T_2 = \inf\left\{ t \geqslant 0 : \hat{X}_s(t) = 0 \text{ or } \sum_{j \neq w} X_j(t) \geqslant N - K \right\}.$$

From the proof of $T > T_1$ above, it follows easily that $T >_s T_2$. Let $V_{ww} = (V_{ww}(t),\ t \geqslant 0)$ denote the portion of V_w destined to leave N_w before visiting s. V_{ww} is a Poisson process of rate $m_w \mu_w \alpha_{ww}$. Similarly define V_{ws}, V_{ss} and V_{sw}. From (3.3), it follows that

$$\lim_{N \to \infty} P\big(V_{ww}(t) + V_{sw}(t) - N^{3/4} \leqslant \hat{A}_s(t) \leqslant V_{ww}(t) + V_{sw}(t) \text{ for all } t \leqslant N^{4/3} \big) = 1,$$

where $(\hat{A}_s(t),\ t \geqslant 0)$ denotes the process counting arrivals into node s in N_w. Thus, one gets the picture that the total customer size in N_w is building up essentially because node s, which started with $N^{4/5}$ customers, is fed by a process which is within $N^{3/4}$ of a Poisson process of rate $m_w \mu_w \alpha_{ws} + m_s \mu_s \alpha_{ss}$. Observe that $\hat{X}_s$ essentially experiences an upward drift of $m_w \mu_w \alpha_{ws} - m_s \mu_s \alpha_{sw}$. From the basic example of [5] (which proves Thm 1.1 for a network consisting of two nodes), it follows immediately that $\lim_{N \to \infty} P(T_2 > ta_N) = 1$ for every $t < 1$, which proves the lower bound in Thm. 3.1.

To prove the upper bound, (3.2), we need to find for each pair of initial conditions for the original network, $x = (n_1, \dots, n_J)$ and $y = (m_1, \dots, m_J)$, having total customer size N, a coupling in time T_{xy} such that

$$\lim_{N \to \infty} \sup_{xy} P(T_{xy} > ta_N) = 0 \quad \text{for every } t > 1. \tag{3.4}$$

The proof of this is parallel to that in [5] with the same sort of modifications in the form of the coupling as those above. For example, the first step is to show that (3.4) is implied by showing

$$\lim_{N \to \infty} \sup_{z} P(T_z > ta_N) = 0 \quad \text{for every } t > 1, \tag{3.5}$$

where z runs over initial conditions for the open network $\boldsymbol{N}_s$, with total customer size at most N, and T_z has the distribution of the time it takes $\boldsymbol{N}_s$ to empty from z. Here $\boldsymbol{N}_s$ denotes the stable open network consisting of the nodes other than node s, and fed by the virtual departure process V_s, of node s. We now indicate how to prove this by modifying the corresponding part of [5].

Given x and y, let $\hat{x}$ and $\hat{y}$ denote the corresponding initial conditions for $\boldsymbol{N}_s$, i.e. just ignore x_s and y_s respectively. We realize $\boldsymbol{N}_s$ on the sample space of the first paragraph, on which we realized $\boldsymbol{N}$. We couple the two closed network processes and the two open network processes in the naive way, i.e. we identify the V_j, and C_j, $1 \leqslant j \leqslant J$, and the routing decisions. We observe that at a time when the open network processes are empty, the closed network processes are also coupled, being in the state where all N customers are at node s. This can be seen by comparing the x process to the $\hat{x}$ process (and the y process to the $\hat{y}$ process) by a coloring idea identical to the one that was outlined above to show $T_1 < T$.

To study the time when both open network processes are simultaneously empty we use a different coloring idea. In each $j \neq s$, start with $\min(x_j, y_j)$ packets colored magenta, $x_j - \min(x_j, y_j)$ packets colored red, and $y_j - \min(x_j, y_j)$ packets colored blue. Observe that initially no node has both red and blue packets. Let the virtual departures from node s (which are arrivals to $\boldsymbol{N}_s$) be colored white. The sum of the white, red and magenta packets is to represent the $\hat{x}$ process, the sum of the white, blue and magenta packets is to represent the $\hat{y}$ process, and the white packets are to represent a version of $\boldsymbol{N}_s$ started empty. We will also arrange things so that no node has both red and blue packets at the same time. Let $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{y}}$, and $\boldsymbol{\emptyset}$ respectively denote the open network process started from $\hat{x}$, started from $\hat{y}$, and started empty. The rules for movement are that if a virtual departure takes place at a node $j \neq s$, and C_j says that this is a real departure for $\hat{\boldsymbol{x}}$, $\hat{\boldsymbol{y}}$, and $\boldsymbol{\emptyset}$, then one moves a white packet. (It is easily seen that there must be at least one white packet at node j in this situation). If C_j says that the departure is real for $\hat{x}$ and $\hat{y}$, but not for $\emptyset$, then one moves a magenta packet. Again, it is easily seen that there must be at least one magenta packet in this situation. If C_j says the departure is real for $\hat{x}$, but not for $\hat{y}$ and $\emptyset$ (resp. real for $\hat{y}$, but not for $\hat{x}$ and $\emptyset$), one moves a red (resp. blue) packet. Once again it is true that whenever this situation occurs a red (resp. blue) packet will be available to be moved. We ensure that there are never both red and blue packets at a node by the rule that if a red (resp. blue) packet moves into a node which has blue (resp. red) packets, we pick a blue (resp. red) packet at the node, the red packet is deleted, and the blue packet is colored magenta.

Let τ_{xy} denote first time at which N_s has only white packets. Let τ_x denote the first time that N_s has only blue and white packets. Note that after this time N_s cannot have any red packets. Also τ_x is stochastically dominated by the time N_s takes to empty from the initial condition $\hat{y}$. Clearly $\tau_{xy} = \max(\tau_x, \tau_y)$. Hence, for any $t > 1$,

$$\begin{aligned} P(\tau_{xy} \leqslant ta_N) &= P(\tau_x \leqslant ta_N \text{ and } \tau_y \leqslant ta_N) \\ &= 1 - P(\tau_x > ta_N \text{ or } \tau_y > ta_N) \\ &\geqslant 1 - \big(P(\tau_x > ta_N) + P(\tau_y > ta_N)\big) \\ &\geqslant 1 - 2\max_z P(T_z > ta_N). \end{aligned}$$

which completes the first step.

The second step is to show that (3.5) is implied by

$$\lim_{N\to\infty} \sup_{j\neq s} P\big(T_j^N > ta_N\big) = 0 \quad \text{for every } t > 1, \tag{3.6}$$

where T_j^N has the distribution of the time it takes N_s to empty when started with N customers at node j and the other nodes empty. To do this, one proves an analog of lemma A.4 of [5], and then follows the proof of lemma 5.2 of [5]. Again the only difference from the proof in [5] is that C_j has to be consulted to determine the color of the packet to be moved when a virtual departure takes place at node j.

Finally one proves (3.6) by defining the set of unstable nodes U_j associated to a node $j \neq s$. Let N_{js} denote the open network consisting of the nodes other than j and s, fed by the virtual departure processes of these nodes. Then we define U_j by saying $i \in U_j$ iff the input rate to i in N_{js} exceeds $m_i\mu_i$. Exactly as in [5], for $k \in U_j$ one has $U_k \subseteq U_j - \{k\}$. It is now easy to prove (3.6) by induction on the cardinality of U_j exactly as in [5]. Intuitively, T_j^N consists of two parts. First, there is the time it takes for node j to empty, at which time the overwhelming portion of the customers left in the network will be in the nodes of U_j. Then one uses inductive hypothesis to estimate the time it takes these customers to leave.

4. A circuit switched network

In this section we consider the transient behaviour of a well studied model for circuit switched networks. This model is originally due to Erlang, see [7], and different aspects of it have been studied by Burman et al., [8], Kelly, [12]–[15], and Ziednis, [23]. A special case of this model has also arisen in the study of multiaccess collision resolution protocols. Lagarias et al., [16].

The network consists of K links. Each link k has C_k circuits. An arriving call must be routed along a fixed route, which is a subset of the set of links. Let the number of possible routes be denoted R. Calls requiring routing along route r arrive according to a Poisson process of rates ν_r. An arriving call along route r demands A_{rk} circuits from each link k along its route, we set $A_{rk} = 0$ if link k is not on route r. If any link along the route of a call does not have enough circuits available to accept the call when it arrives, the call is blocked, i.e. rejected from the system. Each accepted call along route r has a holding time which is exponentially distributed with rate μ_r at the end of which it releases all the circuits it formerly occupied. The arrival processes and the holding times are all mutually independent. See fig. 4.

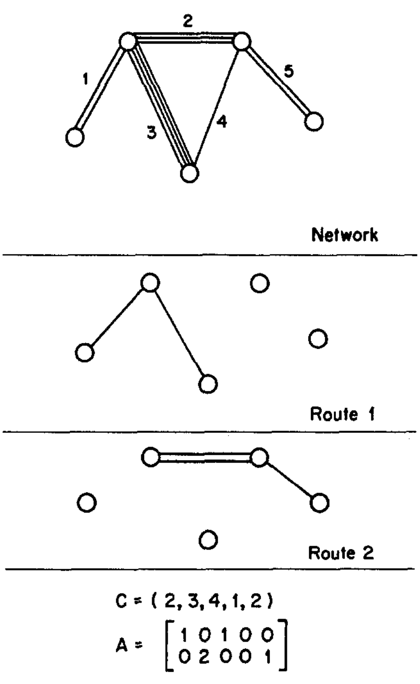

A 5 link circuit switched network with 2 routes

Fig. 4.

Let A denote the $R \times K$ matrix with entries A_{rk}. Let $C = (C_1, \dots, C_K)$. Let X denote the set of vectors $n = (n_1, \dots, n_R)$ such that $nA \leqslant C$. The network can be described by a Markov chain $x = (X(t),\ t \geqslant 0)$ with state space X, the network being in state n when there are n_r calls active along route r. Let ρ_r denote ν_r/μ_r. Then it is easy to see that x admits the stationary distribution π given by

$$\pi(n) = c \prod_{r=1}^{R} \frac{\rho_r^{n_r}}{n_r!}$$

where c is a normalizing constant. To see this, observe that x is the restriction to X of the Markov chain x^∞ on Z_+^R describing the evolution of the above circuit switched network with an infinite number of circuits at each link. It is clear the coordinates of x^∞ evolve as independent M/M/∞ queues with arrival rate ν_r and service rate μ_r respectively. Since x is the restriction of the reversible Markov process x^∞ to a subset X of its state space, its stationary distribution is proportional to the restriction of the stationary distribution of x^∞ to X, see e.g. Kelly, [15], corollary 1.10.

In this section we exhibit a threshold phenomenon in the transient behaviour of this model similar to those described in the earlier sections, as the number of circuits in each link becomes large. The proofs are relatively straightforward and the intuition more superficial. Suppose the network starts in state m, and let

$$d_m^C(t) = \tfrac{1}{2} \sum_{n \in X} | P_m(X(t) = n) - \pi(n) |$$

measure the total variation distance between the distribution of the state at time t and the stationary distribution. Let

$$d^C(t) = \max_{m \in X} d_m^C(t)$$

measure the worst case distance from the stationary distribution over all possible initial conditions. Let

$$a_C = \max_{1 \leqslant r \leqslant R} \left[\log\left(\min_{k \in r} \frac{C_k}{A_{kr}} \right) \frac{1}{\mu_r} \right].$$

Then we have the following

THEOREM 4.1

For every $t < 1$

$$\lim_{C \to \infty} d^C(ta_C) = 1, \tag{4.1}$$

whereas, for every $t > 1$

$$\lim_{C \to \infty} d^C(ta_C) = 0. \tag{4.2}$$

In theorem 4.1, by $C \to \infty$ we mean $\min_{1 \leqslant k \leqslant K} C_k \to \infty$. According to Thm. 4.1, it is natural to consider a_C as the settling time of the circuit switched network. A natural interpretation of this time is that if the network is left in an atypical state, e.g. after a period when some links are down, or after an atypical burst of calls, then one is guaranteed that the network will settle down within time a_C in the precise sense of Thm. 4.1.

As in the preceding sections, the proof consists of two parts. We first prove the lower bound, (4.1), and then the upper bound, (4.2). To prove the lower bound we notice from the form of the stationary distribution that for each $\delta > 0$ it is possible to choose N independent of $\boldsymbol{C}$ such that if Y denotes $\{\boldsymbol{n} \in X: n_r \leqslant N, 1 \leqslant r \leqslant R\}$ then $\pi(Y) > 1 - \delta$ for all sufficiently large $\boldsymbol{C}$ (i.e. if $\min_{1 \leqslant k \leqslant K} C_k$ is sufficiently large). Fix $\boldsymbol{C}$ sufficiently large, and let

$$p = \arg \max_{1 \leqslant r \leqslant R} \left[\log\left(\min_{k \in r} \frac{C_k}{A_{kr}} \right) \frac{1}{\mu_r} \right].$$

Let $\boldsymbol{m} \in \boldsymbol{Z}_+^R$ be given by

$$m_p = \min_{k \in p} \frac{C_k}{A_{kr}}, \qquad m_r = 0, \quad r \neq p.$$

Then $\boldsymbol{m} \in X$. Let T denote the hitting time of Y when the network is started in $\boldsymbol{m}$. From the generalities in section 2, in order to establish the lower bound it suffices to prove that

$$\lim_{C \to \infty} P(T > ta_C) = 1 \quad \text{for each } t < 1.$$

Consider starting the network in the state $\boldsymbol{m}$, but with no arrivals. Let T_1 be the corresponding hitting time of Y. An elementary coupling argument shows that $T >_s T_1$. Let $(S_k, 1 \leqslant k \leqslant m_p)$ be independent exponentially distributed random variables with mean μ_r^{-1}. Let $S_{(m_p)} \leqslant S_{(m_p - 1)} \leqslant \dots \leqslant S_{(1)}$ denote the order statistics of $(S_k, 1 \leqslant k \leqslant m_p)$. Clearly $T_1 =^d S_{(N+1)}$. Thus we have,

$$\begin{aligned} P(T > ta_C) \geqslant P(T_1 > ta_C) &= P\left(S_{(N+1)} > ta_C\right) \\ &= \sum_{m=N+1}^{m_p} \binom{m_p}{m} m_p^{-tm} \left(1 - m_p^{-t}\right)^{m_p - m} \\ &= P\left(B\left(m_p, m_p^{-t}\right) > N + 1\right), \end{aligned}$$

where $B(n, p)$ as usual denotes the sum of n coin tosses with success probability p. A straightforward use of Chebyshev's inequality shows that the right hand side goes to 1 for any $t < 1$, as $\boldsymbol{C}$ (and hence m_p), grows without bound.

To prove the upper bound, let $\boldsymbol{m}$ and $\boldsymbol{n}$ be arbitrary initial conditions for the network. Define $\boldsymbol{v}$ by $v_r = \max(m_r, n_r)$. Let w_r denote $\min_{k \in r} C_k / A_{kr}$, and note that $v_r \leqslant w_r$, $1 \leqslant r \leqslant R$. We construct versions of $\boldsymbol{x}$ started at $\boldsymbol{m}$ and at $\boldsymbol{n}$, and a version of $\boldsymbol{x}^\infty$ started at $\boldsymbol{v}$ on the same sample space in the naive way. That is, our sample space supports R arrival processes, which are Poisson processes of rates ν_r respectively, and arrivals are identified for the three versions. Each arrival brings with it its own service time, which is an exponential random variable of the appropriate mean. To complete the prescription, the sample space supports service times for each of the v_r initial calls of each route r. For each r, m_r, resp. n_r of these are identified with the service times of the initial calls in the respective version of $\boldsymbol{x}$ being constructed.

Clearly the versions of $\boldsymbol{x}$ starting from $\boldsymbol{m}$ and $\boldsymbol{n}$ are coupled at the time that the version of $\boldsymbol{x}^\infty$ started at $\boldsymbol{v}$ empties. But the components of this version of $\boldsymbol{x}^\infty$ evolve as independent M/M/∞ queues with arrival rates ν_r, service rates μ_r and starting at v_r respectively. Let $T_v(\nu, \mu)$ denote the time to empty for an M/M/∞ queue with arrival rate ν and service rate μ and started at v. It remains to study the distribution of this time.

LEMMA 4.2

Let $v \leqslant w$. Then $T_v(\nu, \mu) <_s T_w(\nu, \mu)$.

Proof

Obvious by naive coupling.

LEMMA 4.3

Let V be a Poisson process of rate ν. Let $\boldsymbol{s} = (S_n, n \geqslant 0)$ be an i.i.d. exp(ν) distributed sequence independent of V. We think of S_n as the service time brought in by the nth point of V. Let $(X(t), t < 0)$ denote the queue size process in the corresponding M/M/∞ queue, with $X(0) = 0$.

Let $M_v(\mu)$ be independent of V and $\boldsymbol{s}$, having the distribution of the maximum of v independent exp(μ) distributed random variables. Let

$$T = \inf\{t > M_v(\mu) : X(t) = 0\}.$$

Then $T_v(\nu, \mu) =^d T$.

Proof

We can couple an $M/M/\infty$ queue started with v customers to an M/M/∞ queue started empty by pretending that arrivals are always assigned to servers different from the ones occupied by the initial v customers. The claim is now obvious. □

LEMMA 4.4

Let V, $\boldsymbol{s}$, and $M_v(\mu)$ be as above with the same interpretation. Let $T_\pi(\nu, \mu)$ have the distribution of the time to empty in an M/M/∞ queue started with the stationary distribution, and be independent of $M_v(\mu)$. Then

$$T_v(\nu, \mu) <_s M_v(\mu) + T_\pi(\nu, \mu).$$

Proof

We couple the M/M/∞ process started empty to the M/M/∞ process started with the stationary distribution in the naive way, i.e. identifying arrival times and the service times they bring in with them. We also realize a version of $M_v(\mu)$ on the same sample space which is independent of the M/M/∞ processes. From this it follows that

$$\inf\{t > M_v(\mu) : X(t) = 0\} \leqslant \inf\{t > M_v(\mu) : X^{\pi}(t) = 0\}, \tag{4.3}$$

where $(X^{\pi}(t),\ t \geqslant 0)$ denotes the stationary queue size process. But the right hand side of (4.3) has the distribution of $M_v(\mu) + T_{\pi}(\nu, \mu)$, by stationarity. The claim follows. □

Let $t = 1 + \epsilon > 1$. Let $\boldsymbol{m}$ and $\boldsymbol{n}$ range over admissible initial conditions for fixed $\boldsymbol{C}$. From our coupling and lemmas 4.2–4.4, we have

$$\begin{aligned} \sup_{\boldsymbol{mn}} P(T_{\boldsymbol{mn}} > t a_C) &\leqslant \sum_{r=1}^{R} P\left(T_{w_r}(\nu_r, \mu_r) > t a_C\right) \\ &\leqslant \sum_{r=1}^{R} P\left(M_{w_r}(\mu_r) > \left(1 + \frac{\epsilon}{2}\right) a_C\right) + P\left(T_{\pi_r}(\nu_r, \mu_r) > \frac{\epsilon}{2} a_C\right). \end{aligned} \tag{4.4}$$

But it is easy to see that

$$\lim_{v \to \infty} P\left(M_v(\mu) > \left(1 + \frac{\epsilon}{2}\right) \frac{\log v}{\mu}\right) = 0.$$

Further, we have

$$a_C > \frac{\log w_r}{\mu_r}, \quad 1 \leqslant r \leqslant R$$

and each $w_r \to \infty$ as $\boldsymbol{C} \to \infty$. Since in addition, $T_{\pi}(\nu, \mu)$ is clearly of finite mean, we have that each of the terms on the right hand side of (4.4) goes to zero as $\boldsymbol{C} \to \infty$.

The upper bound now follows easily from the general reasoning in section 2. □

5. A database model with exclusive and non-exclusive locks

In this section we consider the transient behaviour of a model introduced by Mitra, [17], in order to study concurrency control in databases, and in particular to quantify the worth of non-exclusive locks. This model generalizes an earlier database model of Mitra and Weinberger, [18], which can be interpreted as a special case of the model considered in the previous section.

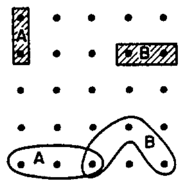

Transaction	Exclusively Locked Items	Non-exclusively Locked Items
A	(1,4) (1,5)	(1,1) (2,1) (3,1)
B	(4,4) (5,4)	(3,1) (4,2) (5,1)

A database with 25 items, shown an admissible state with 2 active transactions

Fig. 5.

The database consists of N items. Transactions require access to subsets of these items. There are p classes of transactions. Each transaction of class σ, $1 \leqslant \sigma \leqslant p$ requires access to $j_\sigma + k_\sigma$ items, requiring exclusive locks on j_σ of these and non-exclusive locks on the other k_σ. See fig. 5. The rationale is that if a transaction needs to read an item, then one can allow the corresponding lock to be nonexclusive. To maintain consistency however, it is often necessary to exclusively lock the items that are being written on. For more on database concurrency control, see e.g. Papadimitriou, [19].

We assume that transaction requests of class σ arrive according to a Poisson process of rate τ_σ. Further, for each class σ, we assume that all possible transactions of this class ar equally likely. Thus the arrival rate of any individual transaction request of class σ is

$$\lambda_\sigma = \frac{\tau_\sigma}{\binom{N}{j_\sigma + k_\sigma}\binom{j_\sigma + k_\sigma}{j_\sigma}}.$$

Let W_d denote the set of exclusively locked items and R_d the set of non-exclusively locked items when a transaction request arrives. Let W_a (resp. R_a) denote the set of items for which the arrival requests exclusive (resp. non-exclusive) locks. The request is accepted iff $W_a \cap (W_d \cup R_d) = \emptyset$ and $R_a \cap W_d = \emptyset$. Otherwise the request is blocked, and is considered as being rejected once and for all (although in practice the request will return to try again later). Once accepted, a transaction of class σ completes after an exponentially distributed time with mean μ_σ^{-1}. The arrival processes and completion times are all mutually independent. Let ρ_σ denote $\lambda_\sigma/\mu_\sigma$.

The system can be modelled as a Markov chain $\boldsymbol{x} = (X(t),\ t \leqslant 0)$, whose state space X consists of *admissible* lists of transactions. Let $\{T_{11}, \ldots, T_{1c(1)}, \ldots, T_{p1}, \ldots, T_{pc(p)}\}$ be a list of transactions. The list is called admissible iff the exclusively locked items of each transaction in the list are pairwise disjoint, and the union of these is disjoint from the union of the non-exclusively locked items. The stationary distribution of the process is given by

$$\pi(x) = c \prod_{\sigma=1}^{p} \rho_\sigma^{c(\sigma)}$$

where x is an admissible list of transactions, $c(\sigma)$ the number of transactions of class σ in x, and c is a normalizing constant. One can see this by verifying detailed balance. The vector $(c(1), \ldots, c(p))$ is called the *concurrency* of x.

Consider this model as the number of items becomes large, with the number of classes of transaction requests and the overall arrival rate of each class remaining fixed. Suppose the database starts in state x and let

$$d_x^N(t) = \tfrac{1}{2} \sum_{y \in X} | P_x(X(t) = y) - \pi(y) |$$

measure the total variation distance between the distribution of the state at time t and the stationary distribution. Let

$$d^N(t) = \max_{x \in X} d_x^N(t)$$

measure the worst case distance from the stationary distribution over all possible initial conditions. Let

$$a_N = \frac{\log N}{\min\limits_{1 \leqslant \sigma \leqslant p} \mu_\sigma}.$$

Then we have the following

THEOREM 5.1

For every $t < 1$

$$\lim_{N \to \infty} d^N(ta_N) = 1, \tag{5.1}$$

whereas, for every $t > 1$

$$\lim_{N \to \infty} d^N(ta_N) = 0. \tag{5.2}$$

According to theorem 5.1, it is natural to consider a_N as the settling time of the database. The interpretation is that if the database is found in an atypical state, perhaps after failures or after a period of atypically large load, one is guaranteed that it will settle to typicality within time a_N.

Proof

We first determine the stationary probability that the state has concurrency $\boldsymbol{c} = (c_1, \ldots, c_p)$. Clearly this is

$$\pi(\boldsymbol{c}) = cF(\boldsymbol{c}) \prod_{\sigma=1}^{p} \rho_\sigma^{c_\sigma},$$

where $F(\boldsymbol{c})$ denotes the cardinality of the set of admissible states with concurrency $\boldsymbol{c}$. It is elementary to see that

$$F(\boldsymbol{c}) = \frac{N!}{\left(N - \sum_{\sigma=1}^{p} c_\sigma j_\sigma\right)! \left(\prod_{\sigma=1}^{p} i_\sigma !\right)\left(\prod_{\sigma=1}^{p} (j_\sigma !)^{i_\sigma}\right)} \prod_{\sigma=1}^{p} \binom{N - \sum_{\sigma=1}^{p} c_\sigma j_\sigma}{k_\sigma}^{i_\sigma}.$$

Further it is reasonably straightforward to see that

$$\lim_{N \to \infty} \pi(\boldsymbol{c}) = \prod_{\sigma=1}^{p} \frac{1}{c_\sigma !} \left(\frac{\tau_\sigma}{\mu_\sigma}\right)^{c_\sigma} \mathrm{e}^{-(\tau_\sigma/\mu_\sigma)}.$$

From this it follows easily that for any $\delta > 0$ it is possible to choose K independent of N such that if Y denotes $\{x \in X\colon c_\sigma(x) \leqslant K,\ 1 \leqslant \sigma \leqslant p\}$, then $\pi(Y) > 1 - \delta$ for all sufficiently large N. Let $\mu_s = \min_{1 \leqslant \sigma \leqslant p} \mu_\sigma$. Let x denote an initial condition where there are

$$\left\lfloor \frac{N - k_s}{j_s} \right\rfloor$$

active transactions of type s and these have precisely the same set of non-exclusive locks. Clearly there is an admissible initial condition of this type. Let T have the distribution of the time to hit Y starting from x. From the generalities in section 2, it suffices to prove that $\lim_{N \to \infty} P(T > ta_N) = 1$ for each $t < 1$ in order to establish the lower bound, (5.1).

Consider starting the database in x but with no arriving requests. Let T_1 be the corresponding hitting time of Y. An elementary coupling argument shows that $T >_s T_1$ and so it suffices to prove that $\lim_{N \to \infty} P(T_1 > ta_N) = 1$ for each $t < 1$. But T_1 has the distribution of the $K + 1$ st largest of a family of

$$\left\lfloor \frac{N - k_s}{j_s} \right\rfloor$$

independent exponentially distributed random variables with mean μ_s^{-1}, so this is immediate, as in the preceding section.

To prove the upper bound, (5.2), let x and y be arbitrary admissible initial conditions, with respective concurrencies $\boldsymbol{c}(x)$ and $\boldsymbol{c}(y)$. For $1 \leqslant \sigma \leqslant p$, let $v_\sigma = \max(c_\sigma(x), c_\sigma(y))$. On the same sample space, we construct versions of x

started at x and y respectively and a family of p independent M/M/∞ queues with arrival rates τ_σ and service rates μ_σ respectively and started at $\boldsymbol{v} = (v_1, \dots, v_p)$. Our sample space supports p Poisson processes of rates τ_σ respectively, these represent the overall arrivals of transaction requests in the first two versions, and arrivals to the respective M/M/∞ queues. It also supports p families of i.i.d. transaction identifying random variables, for the nth transaction request of class σ to arrive, the nth random variable from the σ family tells precisely for which items exclusive and non-exclusive locks are being requested. Each arrival brings with it its own service time, which is an exponential random variable of the appropriate mean. To complete the prescription, the sample space supports service times for each of the v_σ initial customers in each M/M/∞ queue. $c_\sigma(x)$, resp. $c_\sigma(y)$ of these are identified with the service times of the initially active transactions in the respective version of $\boldsymbol{x}$ being constructed.

Clearly the versions of $\boldsymbol{x}$ starting from x and y are coupled at the first time, T_{xy}, that all the M/M/∞ queues are simultaneously empty. But we have

$$v_\sigma \leqslant \left\lceil \frac{N}{j_\sigma} \right\rceil.$$

From lemmas 4.2–4.4 we have, for every $t > 1$,

$$\sup_{xy} P\left(T_{xy} > ta_N\right) \leqslant \sum_{\sigma=1}^{p} P\left(T_{\left\lceil \frac{N}{j_\sigma} \right\rceil}(\tau_\sigma, \mu_\sigma) > ta_N\right). \tag{5.3}$$

Also,

$$a_N \geqslant \frac{\log N}{\mu_\sigma}$$

for each $1 \leqslant \sigma \leqslant p$. It follows easily that the right hand side of (5.3) goes to 0 as $N \to \infty$.

The proof of the upper bound can now be completed by following the general reasoning in section 2. □

6. Concluding remarks

We have examined the transient behaviour of some commonly studied models of communication networks and databases, using techniques similar to those we used in Anantharam, [4], [5], to study the transient behaviour of closed networks of ./M/1 nodes. As such, the proofs in sections 4 and 5 are rather simpler than those in [5], and the threshold phenomenon is less surprising. However, results of this type are valuable in providing insight into the dynamics of real world systems, and one hopes they will have an impact on simulation practice.

References

[1] D. Aldous, Random walks on finite groups and rapidly mixing Markov chains, Preprint, Department of Statistics, University of California, Berkeley, 1987.

[2] D. Aldous and P. Diaconis, Strong uniform times and finite random walks, Advances in Applied Mathematics 8 (1987) 69–97.

[3] D. Aldous and P. Diaconis, Shuffling cards and stopping times, American Mathematical Monthly 93, no. 5 (1986) 333–348.

[4] V. Anantharam, The settling time of a closed Jackson network, *Proc. 25th Allerton Conf. on Communication, Control and Computing*, University of Illinois, Urbana 1987, pp. 684–693.

[5] V. Anantharam, The settling time of a closed Jackson network, Preprint, School of Electrical Engineering, Cornell University, 1988.

[6] P. Bratley, B.L. Fox and L.E. Schrage, *A Guide to Simulation* (Springer Verlag, 1983).

[7] E. Brockmeyer, H.L. Halstrom and A. Jensen, *The Life and Works of A.K. Erlang* (Academy of Technical Sciences, Copenhagen, 1948).

[8] D.Y. Burman, J.P. Lehoczky and Y. Lim, Insensitivity of blocking probabilities in a circuit-switching network, 21 (1984) 850–859.

[9] P. Diaconis and M. Shahshahani, Time to reach stationary in the Bernoulli-Laplace diffusion model, SIAM J. of Math. Analysis 18, no. 1 (1987) 208–218.

[10] W.J. Gordon and G.F. Newell, Closed queueing systems with exponential servers, Operations Research 15 (1967) 254–265.

[11] D. Griffeath, A maximal coupling for Markov chains, Z. Wahr. ver. Geb. 31 (1975) 95–106.

[12] F. Kelly, Stochastic models of computer communication systems, Journal of the Royal Statistical Society, Ser. B 47, No. 3 (1985) 379–395.

[13] F. Kelly, Blocking probabilities in large circuit switched networks, Advances in Applied Probability 18 (1986) 473–505.

[14] F. Kelly, One dimensional circuit switched networks, Annals of Probability 15, No. 1 (1987) 1166–1179.

[15] F. Kelly, *Reversibility and Stochastic Networks* (John Wiley and Sons, Chichester, 1979).

[16] J.C. Lagarias, A.M. Odlyzko and D.B. Zagier, On the capacity of disjointly shared networks, Computer Networks and ISDN Systems 10 (1985) 275–285.

[17] D. Mitra, Probabilistic models and asymptotic results for concurrent processing with exclusive and non-exclusive locks, SIAM Journal of Computing 14, No. 4 (1985) 1030–1051.

[18] D. Mitra and P.J. Weinberger, Probabilistic models of database locking: solutions, computational algorithms and asymptotics, Journal of the ACM 31, No. 4 (1984) 855–878.

[19] C. Papdimitriou, *The Theory of Database Concurrency Control* (Computer Science Press, Maryland, 1986).

[20] S. Ross, *Stochastic Processes* (John Wiley and Sons, New York, 1983).

[21] J. Van der Wal, Monotonicity of throughput in a queueing network, *Abstracts of the Conf. on Queueing Networks and their Applications*, New Brunswick, Jan. 7-9, 1987, p. 127.

[22] J. Walrand, *An Introduction to Queueing Networks* (Prentice Hall Inc., New Jersey, 1987).

[23] I. Ziednis, Quasi-stationary distributions and one-dimensional circuit switched networks, Preprint, Statistical Laboratory, Cambridge University, 1987.