MATHEMATICAL THEORY OF COMMUNICATION NETWORKS

by
Venkat Anantharam

Memorandum No. UCB/ERL M95/22

31 March 1995

MATHEMATICAL THEORY OF COMMUNICATION NETWORKS

by

Venkat Anantharam

Memorandum No. UCB/ERL M95/22

31 March 1995

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

MATHEMATICAL THEORY OF COMMUNICATION NETWORKS

VENKAT ANANTHARAM *
EECS DEPARTMENT
UNIVERSITY OF CALIFORNIA
BERKELEY, CA 94720
ananth@vyasa.eecs.berkeley.edu

Abstract

We describe some recent advances in the mathematical theory of communication networks.

1 Introduction

Developments in telecommunications, manufacturing, and transportation, together with mathematical developments in the theories of interacting particle systems, large deviations, Markov processes, and point processes, have stimulated research on stochastic models having the feature that streams of customers (or packets, or calls) arrive at a system of processing stations, where they occupy resources, move between stations, and eventually leave. Such models are of interest to engineers because, if chosen with sufficient care, they are able to predict the behaviour of engineering systems and aid in their design and operation. They are of interest to mathematicians because they raise a number of fascinating questions, some of which are still unresolved, and also provide a source of examples for more general theories. This area has by now a well established identity, going by the name of *stochastic networks.*

In this article we describe some basic results and sketch some recent advances in this area that have been motivated by problems in telecommunications. The exposition closely follows the pattern of a series of nine fifty-minute lectures delivered by the author at the Fifth Workshop on Stochastic Analysis of Oslo-Silivri held in Silivri, Turkey, in July 1994, at the kind invitation of Professors Hayri Körezlioğlu, Bernt Øksendal, and Süleyman Üstünel. Only a limited range of topics could be covered during the lectures, and the decision of which topics to cover was left to the idiosyncrasies of the author; I would like to apologize in advance to colleagues whose work is not adequately exposed in this article.

*Research supported in part by NSF grant NCR 88-57731

Some references to other survey articles that are useful in developing a broader perspective on stochastic networks are to be found at the appropriate points in the article. The focus throughout is on stochastic models.

This article has been organized into three units, each of them in rough correspondence to a distinct engineering context. In Section 2 we discuss circuit-switched networks, which are useful models for telephony. In Section 3 we discuss datagram networks, which are useful models for the existing generation of data networks. Similar models are also useful in manufacturing and transportation applications. Finally, in Section 4 we discuss more recent questions raised by the drive to merge the telephone and data networks into integrated broadband networks. The nature of the questions that are of interest here has been greatly influenced by the enormous bandwidth available in fiber-optical links.

2 Circuit-switched networks

2.1 The Erlang blocking probability formula

Consider a communication link between two large cities A and B. This link facilitates communication between individuals living in A and those living in B. If the cities are large enough one can plausibly argue, based on the well know limit theorems for a large number of rare events, that the process of call requests between A and B forms a Poisson process, say of rate ν. We assume each accepted call is assigned a fixed amount of bandwidth. We say the link has capacity C circuits if the maximum number of simultaneous calls it can support is C. We assume that a call request arriving when the link is fully loaded is rejected (blocked). If we also make the simplifying assumption that each accepted call holds its assigned bandwidth for an independent exponential duration of mean 1, then we have a simple Markov description of the process of calls in progress. It is a birth and death process on the finite state space $\{0, 1, \ldots, C\}$ with up rates ν and downrate k in state k. Of particular interest is the stationary probability that a call request will be blocked. This is seen to be

$$E(\nu, C) = (\frac{\nu^C}{C!})(\sum_{k=0}^{C} \frac{\nu^k}{k!})^{-1} .$$

This formula is called the *Erlang formula* for blocking probability, see [32].

Two comments should be made at this point. First, the Erlang blocking probability formula gives the time-stationary probability that a call request is blocked. What is more relevant is the stationary probability that an arriving call request finds itself blocked. Here the two probabilities are the same, as a consequence of a property called PASTA (*Poisson Arrivals See Time Averages*), see [80]. The relation between event and time averages is a recurring theme in stochastic networks. For a review of the literature in this area, see Brémaud et al. [16].

The second point is that the Erlang blocking probability formula continues to give the time-stationary probability that a call request is blocked when the assumption of exponential service times is relaxed (for instance, to independent identically distributed service times of mean 1). To determine the stationary blocking probability, all that is important about the service time distribution is its mean. This is an example of an *insensitivity* property, see [18]. Several such insensitivity results are known for stochastic networks. Some of these are described in the text of Walrand [78].

2.2 The circuit switched network model

Consider a graph whose links are numbered $1, \ldots, J$. Link j is assumed to have capacity C_j (a positive integer). In addition there is a finite set of *routes* numbered $1, \ldots, R$. To each route r is associated a J-dimensional column vector of nonnegative integers $a_{jr}, 1 \leq j \leq J$ and a Poisson process of rate ν_r. The interpretation is that call requests along route r arrive at the times of this process and will be accepted iff each link j has at least a_{jr} free circuits, in which case the call holds a_{jr} circuits on link j for an exponentially distributed time of mean 1, after which it simultaneously releases all these resources. The holding times of accepted call requests are independent and independent of the arrival processes. Call requests that arrive to find insufficient resources are rejected (blocked). See Figure 1.

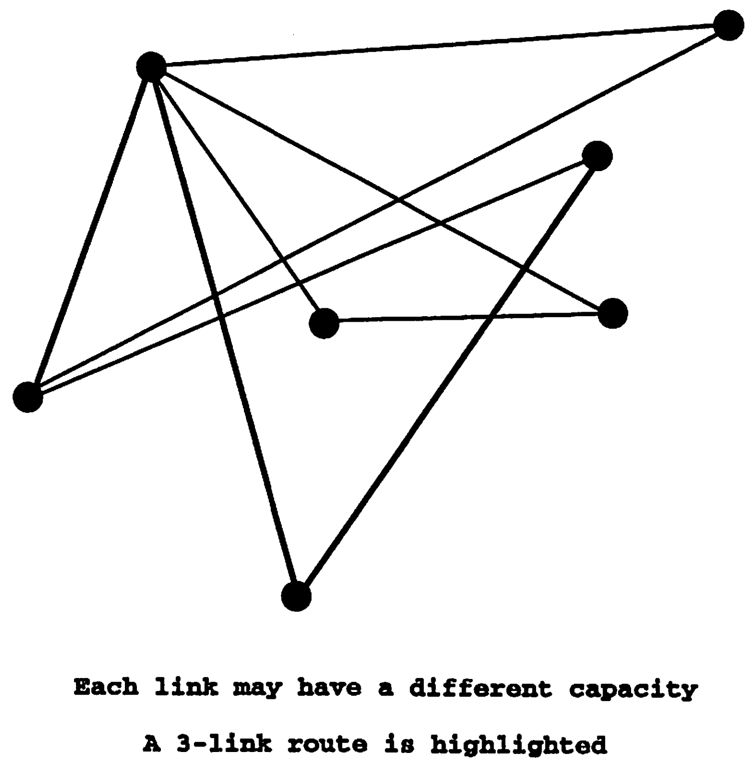

Figure 1: A circuit-switched network with 9 links

Let $n(t) = (n_r(t), 1 \leq r \leq R)$ denote the R-dimensional column vector giving the number of calls along each route that are in progress at time t. This process is a Markov process whose state space is

$$\{n \in \mathbf{Z}_+^R \ : \ An \leq C\}$$

where $A = [a_{jr}]$ is a $J \times R$ matrix, C is the J-dimensional column vector of the C_j, and the inequality is interpreted coordinatewise. This process is time-reversible, and its stationary distribution can be written as

$$\pi(n) = Z^{-1} \prod_r \frac{\nu_r^{n_r}}{n_r!} \tag{1}$$

where

$$Z = \sum_{n \,:\, An \leq C} \prod_r \frac{\nu_r^{n_r}}{n_r!} \,.$$

Let L_r denote the time stationary probability that a call request along route r is blocked. By PASTA, this is the same as the stationary probability that an arriving call request along route r will find itself blocked. Of course the stationary distribution (1) gives us exact formulas for the L_r. Unfortunately these are not of much use in applications, as their use entails computing the normalizing constant Z, which is difficult. In view of this, we seek good approximations for the blocking probabilities.

2.3 The Erlang fixed point approximation

We attempt to define a notion of effective overall arrival rate of requests for individual circuits at each link j, call it ρ_j. Assuming this has been somehow defined, and that the process of overall arrivals is a Poisson process, the Erlang blocking probability formula would say that the stationary probability that link j has all circuits occupied is

$$E_j = E(\rho_j, C_j) \ . \tag{2}$$

Further, the stationary rate at which individual circuits are occupied would be $\rho_j(1 - E_j)$. We now pretend that each circuit request on a link j is rejected independently with probability E_j. This allows us to compute the effective overall rate at which circuits are occupied at link j as

$$\rho_j(1 - E_j) = \sum_r a_{jr}\nu_r \prod_i (1 - E_i)^{a_{ir}} \ . \tag{3}$$

Equations (2) and (3) are to be thought of as a set of fixed point equations for the unknown quantities $E_1, \ldots, E_J$ in terms of the parameters of the model (the matrix A, the capacities, and the arrival rates of call requests). This technique of writing fixed point equations by making independence assumptions is called the *Erlang fixed point approximation technique*. Kelly [45], has proved that there is a *unique* solution to these fixed point equations. The quantity ρ_j is called the *reduced load* at link j. If the assumption that individual circuits block independently approximates reality, one may hope that the approximate equality

$$1 - L_r \approx \prod_j (1 - E_j)^{a_{jr}}$$

is valid in some sense. In [45] it is shown that if one considers a sequence of circuit-switched networks indexed by a paramenter N, with the same matrix A, and with $\frac{1}{N}\nu_r(N)$ and $\frac{1}{N}C_j(N)$ converging to limits as $N \to \infty$ then we have

$$1 - L_r(N) = \prod_j (1 - E_j(N))^{a_{jr}} + o(1) \ .$$

An excellent survey of these and other results on circuit-switched networks is that of Kelly [46].

2.4 Dynamic Routing

The introduction of digital switches and the common-channel signalling system – see for instance [69, Sec. 12-2] for historical perspective – made it possible to consider dynamic routing strategies, where the route of a call between nodes of the network can be chosen adaptively based on traffic conditions. Such strategies are also called *non-hierarchical* routing strategies, in contrast to the hierarchical strategies that were earlier used in the telephone network. Consider the simple network of Figure 2. For each pair of nodes there is Poisson process of call requests of rate ν. If at least one circuit is free on the direct link between the nodes the call is accepted and occupies one circuit for an exponentially distributed time of mean 1. However, if the link is blocked, the call request tries to make the two link connection between the nodes via the remaining node; this is feasible if there is at least one circuit available on each of the other links. If so, the call request is accepted and holds one circuit on each of these links for an exponentially distributed time of mean 1, after which it releases both of them simultaneously.

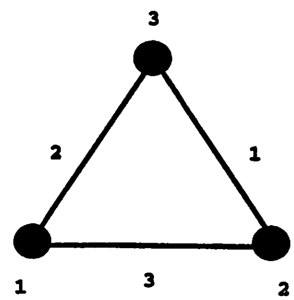

Each link has C circuits

Figure 2: A simple network with dynamic routing

While this network can be described by a finite state Markov process, the state space requires more detail than just the occupancy numbers of the individual links. We are still interested in the stationary probability that an arriving call request is blocked. The Erlang fixed point approach may be adopted to approximate this. Let B denote the stationary probability that an individual link is blocked. The effective arrival rate of requests for circuits on a link is the sum of the direct arrival rate and the arrival rate of requests on each of the other links that have to attempt alternate routing. Assuming that links block independently, this overall rate is seen to be $\nu + 2\nu B(1 - B)$, leading to the fixed point equation

$$B = E(\nu + 2\nu B(1 - B), C) . \tag{4}$$

A sketch of the solutions of this fixed point equation may be found in Figure 1 (i) of Gibbens et al. [37]. Remarkably, for large enough C, there is a range of ν/C where this equation has *multiple solutions.*

The existence of such multiple solutions suggests the possibility of metastable regimes of operation for a circuit-switched network with dynamic routing. In fact, simulations have revealed the existence of hysteresis phenomena in such networks, see [1, 37]. Namely, for certain parameter values, there is more than one qualitatively different regime of operation for the same offered traffic, with the network spending long periods of time in one or the other regime and rapidly moving from one to the other in response to fluctuations in the demand. Intuitively a situation where most calls are using alternate routes is likely to persist for a while because arriving calls will then find the network close to saturation and will be unable to make their direct connections. On the other hand, for the same offered traffic, it might also be the case that if most of the calls in progress are using their direct route, arriving calls will be able to make their direct connection.

This fascinating phenomenon has led to several analytical investigations. A simple model for dynamic routing is considered in [37]. There are n links, each link comprised of C circuits. At each link calls arrive as a Poisson process of rate ν. If the link is not saturated then the call occupies one circuit. If the link is saturated the call chooses two distinct links at random from the $n-1$ remaining links. If neither one is saturated the call occupies one circuit from each of these two links. Otherwise the call is lost. All circuit holding times are exponentially distributed with unit mean, independent of one another and of the arrival times. Further, a call holding circuits from two links is assumed to release them *independently.*

Let $u_k^n(t)$, $0 \le k \le C$ be the fraction of the n links that have k occupied circuits at time t. Then $u^n(t) = (u_0^n(t), u_1^n(t), ..., u_C^n(t))$ evolves on a C-dimensional simplex. It is shown in [37] that as $n \to \infty$, if the initial condition $u^n(0)$ converges weakly to a limit $u(0)$ then the process $(u^n(t), t \in [0,\infty))$ converges weakly to a deterministic process $(u(t), t \in [0,\infty))$ satisfying the following set of differential equations.

$$\begin{aligned}
\dot{u}_0 &= u_1 - (\nu + 2\nu u_C(1-u_C))u_0, \\
\dot{u}_k &= (k+1)u_{k+1} + (\nu + 2\nu u_C(1-u_C))u_{k-1} \\
&- (k + \nu + 2\nu u_C(1-u_C))u_k, \qquad 0 < k < C \\
\dot{u}_C &= -Cu_C + (\nu + 2\nu u_C(1-u_C))u_{C-1}.
\end{aligned} \tag{5}$$

The set of fixed points of this set of equations can be seen to be in one to one correspondence with the solutions of the Erlang fixed point equation (4).

A spatio-temporal version of this model was considered by Anantharam [2]. Let $\mathbf{Z}^d/M$ denote the lattice in $\mathbf{R}^d$ consisting of points all of whose co-ordinates are rational with denominator dividing M. Let W denote $\{0,1,\ldots,C\}$. Let M^* denote $\binom{(2M+1)^d-1}{2}$. Consider a Markov process $(\eta_t^M, t \ge 0)$ on $W^{\mathbf{Z}^d/M}$ which caricatures a circuit switched network with dynamic routing. The Markov process is described by the transitions

$$\begin{aligned}
\eta(x) &\longrightarrow \eta(x) - 1 \text{ at rate } \eta(x)\ , \\
\eta(x) &\longrightarrow \eta(x) + 1 \text{ at rate } \nu \text{ if } \eta(x) \neq C\ , \\
(\eta(x), \eta(y), \eta(z)) &\longrightarrow (\eta(x), \eta(y)+1, \eta(z)+1) \text{ at rate } \nu/M^* \\
&\qquad \text{if } x, y, z \text{ are distinct sites with} \\
&\qquad \eta(x) = C, \eta(y) < C, \eta(z) < C \text{ and } y, z \in x + [-1,1]^d\ .
\end{aligned}$$

Standard techniques, for instance Theorem 3.9 of Liggett [56], ensure that the process is well defined.

In the model each site of the lattice is thought of as representing a portion of a link consisting of C circuits. The value at a site represents the number of occupied circuits in the corresponding link. Thus occupied circuits become free at rate 1 and at each link there is a Poisson process of call requests with rate ν. The dynamic routing is captured by the way call requests are handled : Each call request occupies one circuit on its link if available; if the link is saturated the call randomly picks two other links which are in its $[-1,1]^d$ neighbourhood, and uses one circuit from each of these links if possible. Otherwise the call is blocked and rejected from the system. Note that because we have a compressed lattice, the interaction actually has range M on the scale of links.

For $x \in \mathbf{Z}^d/M$, let $u_M(t,x,k)$ denote $P(\eta_t^M(x)=k)$, $0 \le k \le C$. We may extend the definition of $u_M(t,\cdot,k)$ to $\mathbf{R}^d$ by setting $u_M(t,x,k) = u_M(t,[x]_M,k)$ for $x \in \mathbf{R}^d$, where $[x]_M$ denotes the minimum element in $\mathbf{Z}^d/M$ which dominates x in the usual partial order on $\mathbf{R}^d$. Let $u(0,x,k), 0 \le k \le C$, be continuous functions with bounded derivative and with $\sum_{k=0}^C u(0,x,k) = 1$. Let $u(t,x,k)$, $0 \le k \le C$ denote the solution of the integrodifferential equations

$$\begin{aligned}\frac{\partial u(t,x,0)}{\partial t} &= u(t,x,1) \\ &- \nu\Big(1 + 2^{1-2d}\int\int_{q,r\in[-1,1]^d} u(t,x+q,C)(1-u(t,x+q+r,C))dqdr\Big)u(t,x,0)\,,\end{aligned}$$

$$\begin{aligned}\frac{\partial u(t,x,k)}{\partial t} &= (k+1)u(t,x,k+1) \\ &+\nu\Big(1 + 2^{1-2d}\int\int_{q,r\in[-1,1]^d} u(t,x+q,C)(1-u(t,x+q+r,C))dqdr\Big)u(t,x,k-1) \\ &-\Big(k + \nu\Big(1 + 2^{1-2d}\int\int_{q,r\in[-1,1]^d} u(t,x+q,C)(1-u(t,x+q+r,C))dqdr\Big)\Big)u(t,x,k)\,, \\ &\qquad \text{for } 0 < k < C\,,\end{aligned}$$

$$\begin{aligned}\frac{\partial u(t,x,C)}{\partial t} &= -Cu(t,x,C) \\ &+\nu\Big(1+ 2^{1-2d}\int\int_{q,r\in[-1,1]^d} u(t,x+q,C)(1-u(t,x+q+r,C))dqdr\Big)u(t,x,C-1)\,.\end{aligned}$$

Then we have the following result, see [2].

Theorem 2.1 *Fix $T < \infty$. Suppose that we start $(\eta_t^M)_t$ with initial configuration the product measure having $P(\eta_0^M(x)=k) = u_M(0,x,k)$, $0 \le k \le C$. If $u_M(0,x,k) \to u(0,x,k)$ unformly on compact sets, $0 \le k \le C$, then $u_M(t,x,k) \to u(t,x,k)$ for all $t \in [0,T]$, $x \in \mathbf{R}^d$ and $0 \le k \le C$.*

Theorem 2.1 is a statement about pointwise convergence of probabilities. There is a corresponding functional limit theorem. This functional limit theorem allows us to describe the limit behaviour of an arbitrary choice of spatial integrals $\phi^{(1)}, \ldots, \phi^{(n)}$ at times

$t_1, \ldots, t_n \in [0,T]$ as long as the $\phi^{(i)}$ decrease sufficiently rapidly. This allows, for example, to describe the evolution in time of spatial averages of the state over compact regions of the lattice (which caricatures compact regions of our network with dynamic routing). For details, see [2]. When we look for spatially homogeneous solutions of eqns. (6) which are time invariant, we are led to the same equations as that in the model of [37], so that these are in one to one correspondence with the solutions of the Erlang fixed point equation (4). Thus we see that for large enough C, there is a range of ν over which eqns. (6) admit multiple spatially homogeneous solutions. These may be loosely thought of as different phases associated to the network.

Dynamic routing schemes with symmetry on complete networks were analyzed by Marbukh, [59, 60]. Exploiting the symmetry and making the assumption that in the limit as the number of nodes in the network grow to infinity the evolution of any fixed finite collection of them becomes asymptotically independent (so called *propagation of chaos* assumption; see below) limiting differential equations analogous to (5) were derived for the empirical fraction of links having various occupation numbers. These differential equations also have multiple fixed points for certain ranges of the parameters; the intuitive explanation for this phenomenon is the same as that above. This propagation of chaos assumption was later proved by Graham and Méléard [38]. Graham and Méléard [39] have also proved a fluctuation limit theorem around the law of large numbers that follows from the propagation of chaos proved in [38].

2.5 Propagation of Chaos

Some version of propagation of chaos is involved in all of the above examples, including the writing of Erlang fixed point approximations. What this means is the following : Consider a Markovian system of n identical interacting particles and focus attention on the first p of them. Suppose that as the number of particles increases to infinity the initial distribution of the first p particles becomes asymptotically independent, and the empirical distribution of particle states approaches a limit. In a model having the propagation of chaos property, one can write down a time inhomogeneous Markov process such that the evolution of any given particle is described by this process, started from the appropriate initial distribution. The rates appearing in this process consist of rates associated with autonomous changes of the state of the particle and rates which represent the aggregate effects of interactions of the tagged particle with the large number of other particles. Further, the distinguished finite collection of particles will evolve independently, each according to this time inhomogeneous process. The intuition is that the probability that the distinguished collection of particles interacts with one another becomes asymptotically negligible, and because the interaction is symmetric the interactions with the other particles can be replaced by empirical rates. The terminology comes because the chaotic initial condition (finite collections of particles have asymptotically independent initial conditions) propagates. Propagation of chaos is a well studied property of interacting Markov process models. A recent survey is that of Sznitman [74].

A simple and useful mean field model in which to investigate the hysteresis phenomenon associated to dynamic routing in circuit-switched networks is a model of interacting Markov chains introduced by Uchiyama [76]. Consider a Markov process $X^n(t) = (X_1^n(t), X_2^n(t)...X_n^n(t))$ on $S^n = S \times S... \times S$ where S is a finite set. $(X^n(t))_t$ is characterized by two sets of nonnegative constants $L = \{L^x(y) : x, y \in S,\ x \neq y\}$ and $K = \{K^{x,y}(x', y') : x, y, x', y' \in S,\ (x, y) \neq (x', y')\}$ and an infinitesimal generator given by

$$G_n = \sum_{k=1}^{n} L_k + \frac{1}{n-1} \sum_{1 \leq k < l \leq n} K_{k,l}$$

with

$$L_k \phi(\vec{x}) = \sum_{x'_k \in S} [\phi(\vec{x}'_k) - \phi(\vec{x})] L^{x_k}(x'_k)$$

and

$$K_{k,l} \phi(\vec{x}) = \sum_{x'_k \in S} \sum_{x'_l \in S} [\phi(\vec{x}'_{k,l}) - \phi(\vec{x})] K^{x_k, x_l}(x'_k, x'_l)$$

where $\vec{x} = (x_1, ..., x_n) \in S^n$, ϕ is a real function on S^n, $\vec{x}'_k$ [resp. $\vec{x}'_{k,l}$] is an element of S^n obtained from $\vec{x}$ by replacing x_k [resp. x_k and x_l] with x'_k [resp. x'_k and x'_l], and the sum $\sum_{k<l}$ is taken over all pairs (k, l) such that $1 \leq k < l \leq n$. Also assume that $K^{x,y}(x', y') = K^{y,x}(y', x')$. (More generally a fixed number of exchangeable multi-particle interactions can be considered, instead of just considering two-particle interactions as above.)

We think of S as the state space of an individual particle. Hence $X_k^n(t)$ is the physical state of the kth particle in a system of n identical particles. The process evolves as follows: each particle k evolves autonomously through a Markovian motion according to L_k. Pairwise interaction between particles k and l is controlled by $K_{k,l}$: two particles in state x and y change simultaneously to states x' and y' respectively at rate $\frac{K^{x,y}(x',y')}{n-1}$. The factor $n-1$ is there so that the interaction rate per particle is constant as $n \to \infty$. Another way to think of the pairwise interaction is as follows : for each $y, x', y' \in S$, each particle in state x chooses another particle at random from the remaining $n-1$ at rate $\frac{1}{2} K^{x,y}(x', y')$ and if this particle is in state y they change together to states x' and y' respectively. If the chosen particle is not in state y nothing happens.

Now let $u_i^n(t) = \frac{1}{n} \sum_{l=1}^{n} 1(X_l^n(t) = i)$. Then $u^n(t) = (u_1^n(t), ..., u_{|S|}^n(t))$ is a Markov chain on the $|S|-$dimensional simplex Δ given by

$$\Delta = \{\vec{u} \in R^{|S|}, \sum_{i=1}^{|S|} u_i = 1\}$$

For $i \neq j, 1 \leq i, j \leq |S|$, let T_{ij} be an operator defined on $\vec{u} \in \Delta$ by

$$T_{ij} \vec{u} = \vec{u} + \frac{1}{n}(e_j - e_i)$$

where e_i is the unit vector in the ith direction. Then the infinitesimal generator A_n of the Markov Process $u^n(t)$ on Δ is the operator given by

$$A_n\phi(\vec{u}) = \sum_{\substack{i,j\in S\\ i\neq j}} [\phi(T_{ij}\vec{u}) - \phi(\vec{u})]L^i(j)u_i n \tag{6}$$

$$+ \sum_{\substack{i,j,i',j'\in S\\ i\neq i',(i,i')\neq(j,j')}} [\phi(T_{ij}T_{i'j'}\vec{u}) - \phi(\vec{u})]\frac{K^{i,i'}(j,j')}{n-1}\frac{u_i n u_{i'} n}{2} \tag{7}$$

$$+ \sum_{\substack{i,j,j'\in S\\ (i,i)\neq(j,j')}} [\phi(T_{ij}T_{ij'}\vec{u}) - \phi(\vec{u})]\frac{K^{i,i}(j,j')}{n-1}\frac{u_i n(u_i n-1)}{2} \tag{8}$$

where ϕ is a continuous function on Δ.

Let $u(t) \in \Delta$ evolve according to the following equation started with $u(0)$.

$$\begin{aligned} \dot{u}_i(t) &= \sum_{\substack{j\in S\\ j\neq i}} L^j(i)u_j - \sum_{\substack{j\in S\\ j\neq i}} L^i(j)u_i \\ &+ \sum_{i'\in S}\sum_{j\in S} K_1^{i',j}(i)u_{i'}u_j - \sum_{i'\in S}\sum_{j\in S} K_1^{i,j}(i')u_i u_j \end{aligned} \tag{9}$$

for $i \in S$ where $K_1^{x,y}(x') = \sum_{y'\in S} K^{x,y}(x',y')$.

Then the idea of propagation of chaos is captured by the following theorem.

Theorem 2.2 *For the system of n interacting particles above let $(X_1(t),...,X_p(t))$ denote the state of the first p particles and $u_x^{p+1,n}(t)$ the fraction of particles $p+1 \leq l \leq n$ that are in state x at time t. Let $u^{p+1,n}(t) = (u_x^{p+1,n}(t), x \in S)$. Suppose $(X_1(0),...,X_p(0),u^{p+1,n}(0))$ converges weakly to a product distribution $\mu^{(1)} \otimes ... \otimes \mu^{(p)} \otimes \delta_{u(0)}$ in $E = S^p \times \Delta$.*

Let $u(t)$ solve the ODE (9) starting at $u(0)$. For $1 \leq l \leq p$ let $P^{\mu^{(l)}}$ be the probability measure on $D_E[0,\infty)$ corresponding to the time inhomogeneous Markov chain $X(t)$ with state space S, with initial distribution $\mu^{(l)}$, and such that the rate of jumping from state s to state s' is

$$\lambda(u(t),s,s') = L^s(s') + \sum_{\substack{i,j\in S\\ (i,s)\neq(j,s')}} K^{i,s}(j,s')u_i(t)$$

Then the process $(X_1(t),...,X_p(t),u^{p+1,n}(t))$ converges weakly to a product distribution $P^{\mu^{(1)}} \otimes ... \otimes P^{\mu^{(p)}} \otimes \delta_{u(t)}$ in $D_E[0,\infty)$.

This theorem is due to Uchiyama [76]. A proof of this theorem using modern tools from the theory of weak convergence of Markov processes, see e.g. Ethier and Kurtz [33], can be found in Anantharam and Benchekroun [6], together with applications to computing approximations to sojourn times in networks of queues. The model of [37] can be verified to be a model of this type. In this model Theorem 2.2 gives the additional information that the evolution of any finite collection of links is asymptotically independent in the limit as the total number of links approaches infinity, and that each link evolves according to a time-inhomogeneous birth and death process with up and down rates given explicitly in term of the overall time varying empirical occupation distribution, which follows eqn. (5).

2.6 Large deviations

It is of particular importance to give simple rules which can predict which of the metastable regimes is likely to dominate for a given set of parameter values. One approach to this problem would be via the theory of large deviations. Intuitively, there is a "potential well" associated to each equilibrium, and the equilibrium that is likely to dominate is the one whose associated well is deepest, in that it takes the longest time to escape from the well via a rare fluctuation. See Wentzell and Freidlin [79] for a rigorous formalization of such intuition.

There does not appear to be a solution yet to this challenging problem. To close the section, we briefly describe a recent result that is able to answer a related, albeit much more special question. Consider the model of [76] (with a fixed number of exchangeable multiple particle interactions) in the special case where $S = \{0, 1\}$, i.e. when each of the interacting chains is a 2-state chain. For each n the process $(X^n(t))_t$ is a finite state Markov process, and therefore admits a unique equilibrium distribution. Let α_n denote the stationary distribution of the empirical distribution of the interacting chains. Since $S = \{0, 1\}$ this is a distribution on the set $\{\frac{k}{n}, 0 \leq k \leq n\}$. We may write

$$\alpha_n(\frac{k}{n}) = \binom{n}{k} \exp(n h_n(\frac{k}{n}))$$

for some function h_n. We extend the definition of h_n to $[0,1]$ by linear interpolation. In Anantharam [5], we prove the following result.

Theorem 2.3 *The functions h_n converge uniformly to a limit h. A consequence is that the distributions α_n obey a large deviations principle with action functional*

$$I(u) = -h(u) + D(u, \frac{1}{2}) + p \qquad 0 \leq u \leq 1$$

where

$$D(u, \frac{1}{2}) = u \log 2u + (1-u) \log(2(1-u))$$

and

$$p = \sup_{u \in [0,1]} [h(u) - D(u, \frac{1}{2})] \ .$$

It is of particular interest that the action functional is in general nonconvex. Its local minima correspond to the fixed points of the limiting differential equation (9). The global minima of the action functional may then be considered as corresponding to the dominating regimes for the given parameter values. Arriving at a theorem of this sort for the case of general S in the model of [76], and therefore in the model of [37], would be of great interest.

3 Datagram networks

Data networks work by breaking messages into packets, which are routed through the network and reassembled at the destination. The mathematical analysis and design of such networks involves studying networks of queues.

3.1 The M/M/1 queue

A basic queueing model is the M/M/1 queue. Packets enter a buffer of infinite capacity at the times of a Poisson process of rate λ. Each packet brings in an amount of work which is an exponential random variable of mean μ^{-1}, these variables being independent from packet to packet, and independent of the arrival process. There is a server performing work at rate 1, that operates as follows : on finishing serving a packet it picks an arbitrary packet from the buffer, if any, and proceeds to serve that packet till all the work it has brought in is complete. Packets that have had their work completely served immediately depart the buffer. The first M is a mnemonic for the Poisson (memoryless) character of the arrival process, the second M for the exponential (memoryless) character of the work distributions, the 1 reminds us there is a single server at work. (The queueing literature is littered with a taxonomist's paradise of such abbreviations.) The work brought in by a packet may be taken as representative of its length, or more generally of the time required to carry out some processing of it.

It is conventional to think of the packets as being served in the first-come-first-served (FCFS) order, but as long as the identity of the individual packets is of no concern the exponential assumptions ensure that the order in which packets are served is irrelevant.

Let $(X(t), t \geq 0)$ be the process of total of number of packets in the buffer, including the packet being served, if any. Then $(X(t), t \geq 0)$ is a birth and death process on the non-negative integers $\mathbf{Z}_+$ with up-rate λ and down-rate μ. It admits a stationary distribution if and only if $\lambda < \mu$, in which case, with $\rho = \lambda\mu^{-1}$ denoting the *traffic intensity*, the stationary distribution is

$$\pi(n) = \rho^n(1-\rho) \quad , \; n \geq 0 \; .$$

3.2 The Jackson network model

Let us make three observations. The first is *Burke's theorem*, also called the *output theorem* for the the M/M/1 queue, see [17]. Consider an M/M/1 queue with $\lambda < \mu$ in stationarity as a process defined for all $t \in \mathbf{R}$. Let $(A_t, t \in \mathbf{R})$ denote the (Poisson) arrival process, and $(D_t, t \in \mathbf{R})$ the departure process of packets leaving the queue. Burke's theorem states that $(D_t, t \in \mathbf{R})$ is a Poisson process of rate λ whose past at any time $t \in \mathbf{R}$ is independent of the present state of the buffer, i.e.,

$$(D_s, s \leq t) \amalg X_t \quad , \; t \in \mathbf{R} \; .$$

This is a direct consequence of the time reversibility of a stationary birth and death process. Indeed, the departure process of the forward-time process is precisely the arrival process of the reverse-time process. Nevertheless it is a striking, and even counterintuitive, result as naive intuition would suggest, for instance, that an enormous number of departures just prior to a given time t would result in an increased likelihood of the buffer being empty at time t.

The second observation is that Bernoulli sampling of a Poisson process results in independent Poisson processes. More precisely, if at each time of a Poisson process $(N_t, t \in \mathbf{R})$ of rate λ we draw an independent sample of a Bernoulli random variable whose probability of being 1 is p, and split the process into two streams $(N_t^1, t \in \mathbf{R})$ and $(N_t^0, t \in \mathbf{R})$ consisting respectively of those points for which the value of the Bernoulli random variable was 1 or 0, then these are independent Poisson streams of rates λp and $\lambda(1-p)$ respectively.

The third observation is that the sum of independent Poisson processes is a Poisson process.

We now describe a model of interconnected queues due to Jackson [44] by which to describe a system into which packets enter, move between servers, and eventually leave. A Jackson network has J infinite buffers. Packets arrive from the external world at the times of a Poisson process of rate γ. The outside world is conventionally indexed by 0. An arriving packet is routed to buffer j with probability r_{0j}, with $\sum_{j=1}^{J} r_{0j} = 1$, where it queues in FCFS order. The work required by a packet at node j is an exponential random variable of mean μ_j^{-1}. Each buffer is served by its own server, that works on the leading packet in the buffer at rate 1 till it completes the work required by this packet, after which it immediately begins work on the next packet in the buffer, if any. With probability r_{jk} a packet completing service at node j is routed to buffer k where it queues in FCFS order, and it leaves the system with probability r_{j0}, where $\sum_{k=0}^{J} r_{jk} = 1$. Service times are i.i.d and independent of the arrival process. Routing is Bernoulli, independent of the arrival process and the service times.

We assume that the Jackson network is irreducible, i.e. it is possible for an exogenous arrival to visit any queue before leaving the system. We also assume that it is stable, namely that the solutions of the flow balance equations:

$$\lambda_i = \gamma r_{0i} + \sum_{j=1}^{J} \lambda_j r_{ji}, \qquad 1 \le i \le J \tag{10}$$

satisfy

$$\lambda_i < \mu_i, \qquad 1 \le i \le J \, . \tag{11}$$

Let $X_j(t)$ denote the number of packets in buffer j at time t. and let $X(t) = (X_1(t), \ldots, X_J(t))$. Then $(X(t))_t$ is a Markov process. Note that as long as the focus is on this Markov process, so that the identities of the individual packets are ignored, it is not important that the service at the individual buffers be FCFS, or even that it be *non-preemptive* (i.e. the sever could leave a partially worked-on packet and move to begin work on another packet). All that is important is that the servers be *work-conserving*, i.e.

that they work whenever there is at least one packet in the corresponding buffer. Under the stability and irreducibility conditions the Markov process $(X(t))_t$ admits a unique stationary distribution. The importance of the Jackson network model in applications stems from the simple form of its stationary distribution, which is given by

$$\pi(x_1, \ldots, x_J) = \prod_{i \in [J]} \rho_i^{x_i}(1 - \rho_i) \tag{12}$$

where $\rho_i = \lambda_i \mu_i^{-1}$ is called the load factor at node i. Such a stationary distribution is said to be of *product-form*, since the individual stationary queue sizes are independent. As might be imagined, this facilitates the computation of stationary quantities. Note that it is far from the case that the evolutions of the individual buffer sizes are independent in stationarity. An interesting line of research in stochastic networks is to develop more sophisticated network models with product-form stationary distributions. A recent survey of some of this literature is that of Nelson [63].

Walrand [77] has provided a clever explanation for the product-form stationary distribution. Suppose that we introduce a small delay in the feedback of packets after they have finished service at a buffer and before they are routed to another buffer in the system. Schematically, the situation can be represented as in Figure 3.

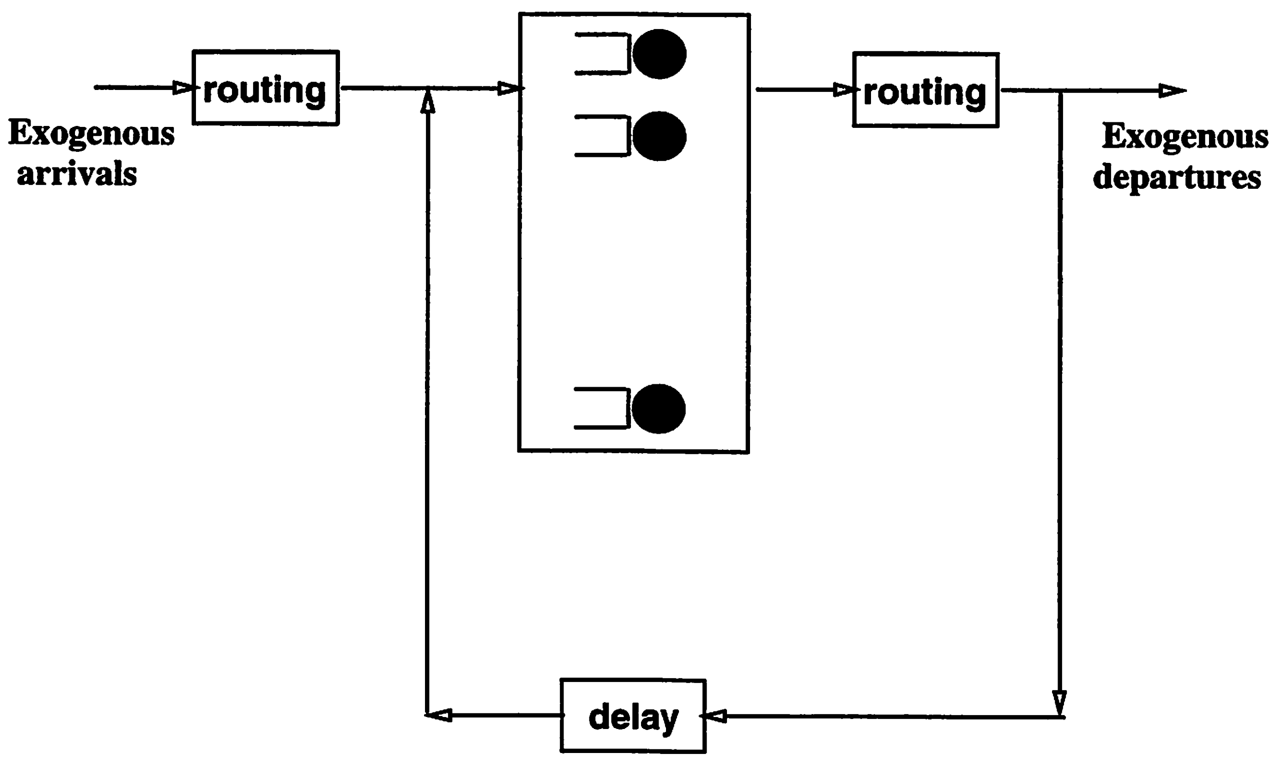

Figure 3: Explaining the product form stationary distribution

Let $(X^\Delta(t))_t$ be the process of buffer sizes in this modified system, where $X^\Delta(t) = (X_1^\Delta(t), \ldots, X_J^\Delta(t))$. We can guess the stationary distribution of $(X^\Delta(t))_t$ using the intuition gained from the M/M/1 queue. Suppose that $X_j^{\Delta}(0-)$ are independent geometric random variables with parameter $\rho_j = \lambda_j \mu_j^{-1}$, $1 \leq j \leq J$ (where λ_j was defined in eqn. (10)). Also assume that the delay line holds independent pieces of Poisson processes of rate $\sum_i \lambda_i r_{ij}$ respectively destined for buffer j, and that these pieces are independent of $X^{\Delta}(0-)$. This sets up the initial conditions at time 0. We now argue that the situation at time Δ is exactly the same. Over the time interval $[0, \Delta)$ the individual buffers are being fed by independent Poisson process of rate λ_j respectively, since this is the result of summing independent Poisson processes of rates γr_{0j} and $\sum_i \lambda_i r_{ij}$ (see eqn. (10)). This process is independent of the initial condition $X_j^{\Delta}(0-)$, which is the stationary initial condition of an M/M/1 queue of arrival rate λ_j and service time of mean μ_j^{-1}, so it follows that $X^{\Delta}(\Delta-)$ is also a collection of independent geometric random variables of parameter ρ_j respectively. But Burke's theorem tells us that the departures from the buffers over the interval $[0, \Delta)$ are independent of this vector ! Since Bernoulli sampling of each Poisson process results in independent Poisson processes, and the sums of independent Poisson processes are Poisson, the delay line once again holds pieces of Poisson processes of rates $\sum_i \lambda_i r_{ij}$ respectively destined for buffer j, that are independent of $X^{\Delta}(\Delta-)$. When we let $\Delta \to 0$ we can understand why the stationary distribution of the Jackson network is product-form.

3.3 Stability of Jackson-type networks

The stability condition (11) for the Jackson network is easy to understand : the effective rate at which work enters a buffer should be strictly less the service rate. Of course, discussing "effective service rate" presupposes that the network is stable. One of the main concerns in stochastic networks over recent years has been to come to grips with the question of necessary and sufficient stability conditions in more general queueing network models.

Consider the following generalization of the Jackson network model : We retain the Bernoulli routing feature, but generalize the process of packet arrival times to a general renewal process of rate γ, and the service times of packets at the individual buffers to general independent identically distributed service times with mean service time μ_j^{-1} at buffer j. We also insist that service at the individual buffers is FCFS, non-preemptive, and work-conserving. Intuition would suggest that if the flow balance equations (10) satisfy (11), the resulting process is stable in some sense.

This problem turned out to be very difficult to resolve. Early works on this problem are due to Borovkov [12], and Sigman [73]. Relatively satisfactory solutions to the problem have only been found quite recently, see Foss [34, 35], Meyn and Down [61], Chang et al. [20], Baccelli and Foss [10], and Dai [25].

We now give a glimpse of the elegant solution of this problem in [10]. Here the problem of stability of datagram networks is approached from a very general point of view. Consider first a broad generalization of the single server M/M/1 queue. On a sample space $(\Omega, \mathcal{F}, \mathcal{P})$ admitting a shift θ under which P is ergodic, we are given nonnegative random variables (σ_0, τ_0) satisfying $E[\sigma_0] = \mu^{-1} < \infty$ and $E[\tau_0] = \lambda^{-1} < \infty$. Let $(\sigma_n, \tau_n) = (\sigma_0 \circ \theta^n, \tau_0 \circ \theta^n)$. Thus $\{\sigma_n, \tau_n\}_n$ is a stationary ergodic sequence. σ_n has the interpretation of the work brought in by the nth customer to a server and τ_n of the interarrival time between the arrival of the nth customer and the $n+1$st customer. The server works at rate 1 if there is work in the system. Let W_n denote the workload in the system seen by the nth customer. Then, starting from some initial condition, the workload evolves according to the equation

$$W_{n+1} = (W_n + \sigma_n - \tau_n)^+ \tag{13}$$

This equation is called the *Lindley equation.* We ask for what parameter values this recursion admits a stationary solution, i.e., a proper random variable W_0 such that, with $W_n = W_0 \circ \theta^n$, we have $(W_n)_n$ satisfying (13). The M/M/1 queue is the special case where the sequence $\{\sigma_n, \tau_n\}_n$ is i.i.d. with σ_0 and τ_0 independent exponential random variables. For the M/M/1 queue, we know that the condition for existence of a stationary solution to (13) is precisely $\lambda < \mu$. The following remarkable result is due to Loynes [57].

Theorem 3.1 *If $\lambda < \mu$ the Lindley equation admits a unique stationary solution. If $\lambda > \mu$ there is no stationary solution to the Lindley equation.*

Proof:

The proof is representative of a large number of results of this type, so it is worthwhile to sketch the ideas. Fix $m \geq 0$. We define random variables $(W_n^m, n \geq -m)$ with $W_{-m}^m = 0$ and obeying Lindley's equation

$$W_{n+1}^m = (W_n^m + \sigma_n - \tau_n)^+ \ . \tag{14}$$

We observe that W_{n+1} is a monotone increasing function of W_n in (13). Thus $(W_n^m, m \geq 0)$ is increasing in m for fixed n, and has a limit W_n^∞ which obeys

$$W_{n+1}^\infty = (W_n^\infty + \sigma_n - \tau_n)^+ \ . \tag{15}$$

Since $W_{n+1}^m = W_n^{m+1} \circ \theta$, we have $W_{n+1}^\infty = W_n^\infty \circ \theta$. It remains to consider when W_n^∞ is proper. We may write

$$W_n^{m+1} \circ \theta = W_{n+1}^m = W_n^m - W_n^m \wedge (\tau_n - \sigma_n) \ . \tag{16}$$

Taking expectations, we get

$$E[W_n^m \wedge (\tau_n - \sigma_n)] \leq 0 \tag{17}$$

and so

$$E[W_n^\infty \wedge (\tau_n - \sigma_n)] \leq 0 \ . \tag{18}$$

Since $(W_n^\infty)_n$ is a stationary ergodic sequence, $P(W_n^\infty = \infty)$ is either 0 or 1. From (18) we see that $P(W_n^\infty = \infty) = 1$ implies that $E[\tau_n] \leq E[\sigma_n]$, i.e., $\lambda \geq \mu$. This proves that if $\lambda < \mu$ then a stationary solution to the Lindley equation exists. For the situation when $\lambda > \mu$, first note that $(W_n^\infty)_n$ constructed above is the minimal θ-invariant solution of (13), and that

$$W_0^\infty = (\sup_n \sum_{k=1}^{n} (\sigma_{-k} - \tau_{-k}))^+ = \infty \tag{19}$$

when $\lambda > \mu$. For more details regarding the solutions of (13) see Baccelli and Brémaud [9, Sec. 2.2]. □

The stability question for datagram networks can be studied in the stationary ergodic framework. A fundamental distinction that emerges is between a station-centered and a customer-centered point of view, see [10]. In the station-centered model, called the *Jackson-type network*, the service times and routing variables are associated to the stations, and handed out to the packets by the server as they are picked up for service. One can visualize identical, featureless packets moving around the network and queuing up in FCFS fashion at the individual nodes. When it is the turn of such a packet to receive service at a node it picks up its service time from a list maintained at the node; when this service is complete it then picks up a routing variable from a list maintained at the node to decide which node to move to next. In contrast to this, in the customer-centered model, called the *Kelly-type network*, the service and routing variables are associated to the individual packets. One can visualize the packets arriving with marks giving their entire route through the network, together with the service times the packet will require at each visit to each node along the route. As before the service at each node is in FCFS order of arrival, but the mark is carried by the packet throughout its route. We discuss each of these models in turn.

The generalized Jackson network with renewal arrivals, i.i.d. service, and Bernoulli routing, can be considered as either a Jackson-type network or a Kelly-type network. Baccelli and Foss [10] have given a very satisfactory solution to the stability problem for Jackson-type networks. Consider a network consisting of a fixed finite number of stations, each of which has infinite waiting room and a single server that works at unit rate. Consider a stationary ergodic marked point process of arrivals, each of which brings with it a route through the nodes of the network and a service variable for each visit to each node along the route. This process can be described on a sample space $(\Omega, \mathcal{F}, \mathcal{P})$ admitting a shift θ under which P is ergodic, by giving a pair (ξ_0, τ_0) : the variable ξ_0 describes the route through the nodes of the network and the corresponding service variables for each visit to each node along the route, and τ_0 gives the time to the next arrival. We now visualize the entire mark peeling off the arriving customer at the time of its arrival and joining lists of service and routing variables maintained at the individual *nodes*, in sequence, at the tail of such lists. The featureless packets now move around the network picking up service and rounting variables that are handed to them from these lists by the *nodes*.

A key observation is that this mechanism of peeling off the marks of the arriving packets and attaching them to lists at the appropriate nodes ensures that service and routing variables are always available at a node when needed. What is meant by this statement is the following : Consider an initially empty network (also with empty service time and routing lists at each node) and suppose that n packets arrive at the network at times $t(1) < t(2) < \ldots < t(n)$. Each packet brings a mark (a route through the network and a service time requirement at each visit to each node along the route). This mark is peeled off immediately on arrival and distributed among the appropriate lists at the appropriate nodes. The now featureless packets move through the network, queueing at the nodes in FCFS order and picking up service time and routing variables from the lists. Then we will *not* have a situation where a packet looking for a service time variable to pick up to enter service at a node or a routing variable to pick up on finishing service at a node finds the list empty. This statement needs a proof; indeed, with the mechanism of peeling off marks being used, it is possible for a packet to use a service time or routing variable brought in by a packet that arrives *after* it does.

A second key observation is a monotonicity property proved by Foss [34], and Shantikumar and Yao [71]; see also [20, Prop. 4.1] and [10], Theorem 10. This says that, with the mechanism of peeling off marks on arrival that was just described, when we consider the situation with n customers entering into an empty network with empty lists, then delaying the arrival times or increasing any of the service requirements delays the times at which service completions take place; in fact, for any nodes k and l and any $j \geq 1$, the jth service at node k that sends a packet to node l will be delayed. Further the time by which these services are delayed can be bounded in terms of the time by which the arrivals are delayed and the increases in the individual service times at the nodes, see [10, Corollary 3].

These observations now allow us to set into motion a machinery to use the Loynes's scheme and exploit Kingman's subadditive theorem by considering the effect of dilations and time shifts of the arrival process. The essential features of this process have been abstracted by Baccelli and Foss [11], and results in a *saturation rule* for stability. Namely if we consider a limiting system in which the network is started empty and immediately receives an infinite number of arrivals, and compute the rate at which packets are released from this scenario then the original network will admit a stationary regime for all arrival rates strictly less than this saturated evacuation rate, and will not admit a stationary regime if the arrival rate strictly exceeds this saturated evacuation rate. In [10] it is shown that this characterization of stability is precisely that given by the requirement (11) on the λ_i satisfying (10). The construction of the stationary regime for the network when (11) holds follows the lines of Loynes scheme, because the monotonicity property proved above allows the variables of interest in the network to be described in terms of a monotone stochastic recursion.

3.4 Weak Solutions of Stochastic Recursions

The recursion (13) is an example of a *stochastic recursion*. Let $(\Omega, \mathcal{F}, \mathcal{P})$ be a probability space admitting a shift θ under which P is stationary and ergodic. Let $(E, \mathcal{E})$ be a Polish space and φ_0 a random variable defined on $(\Omega, \mathcal{F}, \mathcal{P})$ that takes values in the space of measurable maps from $(E, \mathcal{E})$ into itself. Let $\varphi_n(\omega) = \varphi_0(\theta^n \omega)$. Then, under P, $\{\varphi_n, n \geq 0\}$ is a stationary ergodic sequence of random maps from $(E, \mathcal{E})$ into itself. In many applications the stability question can be framed as one of finding the conditions under which recursions of the form

$$x_{n+1} = \varphi_n(x_n) \tag{20}$$

admit a stationary solution. For instance (13) is an equation of this type with $(E, \mathcal{E}) = (\mathbf{R}, \mathcal{B})$ and $\phi_0(x) = (x + \sigma_0 - \tau_0)^+$. Note that ϕ_0 is monotone. By and large, all such recursions that have been successfully studied in the literature exploit some kind of monotonicy of ϕ_0, and are handled by proceeding pathwise using a version of the Loynes' scheme outlined in the proof of Theorem 3.1. Recently Anantharam and Konstantopoulos [8], have developed another approach to solving such stochastic recursions, which does not require monotonicity, but is based on the weakening of the solution concept, along lines also proposed in [15]. We briefly sketch this approach; for details, see [8].

A *weak solution* of the recursion (20) is a measure Q on a measurable space $(\tilde{\Omega}, \tilde{\mathcal{F}})$ admitting a measurable shift $\tilde{\theta}$ such that Q is $\tilde{\theta}$-invariant, and a pair of random variables X_0 and Φ_0 on $(\tilde{\Omega}, \tilde{\mathcal{F}})$ taking values respectively in E and in the space of measurable maps from $(E, \mathcal{E})$ into itself, such that

$$X_0 \circ \tilde{\theta} = \Phi_0(X_0) . \tag{21}$$

In [8] we observe that it is often possible to construct such a weak solution along the lines of the skew product construction in ergodic theory (see, e.g. Krengel [50]), even in the absence of any kind of monotonicity of ϕ_0. Assume that $(\Omega, \mathcal{F})$ is a Polish space. Consider the product space $\Omega \times E$, endowed with the product σ-field $\mathcal{F} \otimes \mathcal{E}$ and the new measurable shift $\Theta(\omega, x) : \Omega \times E \to \Omega \times E$ defined by

$$\Theta(\omega, x) = (\theta\omega,\ \varphi_0(\omega)[x]). \tag{22}$$

Note the following composition rule.

$$\begin{aligned} \Theta^n(\omega, x) &= (\theta^n\omega,\ \varphi_0(\theta^{n-1}\omega)\varphi_0(\theta^{n-2}\omega)\ldots\varphi_0(\omega)[x]) \\ &= (\theta^n\omega,\ \varphi_{n-1}(\omega)\varphi_{n-2}(\omega)\ldots\varphi_0(\omega)[x]) , \\ \Theta^{n+m}(\omega, x) &= \Theta^n(\Theta^m(\omega, x)) . \end{aligned}$$

In [8], the following result is proved.

Theorem 3.2 *Let Q_0 be a probability distribution on $\Omega \times E$ whose Ω marginal is P. Let Q_n denote the probability distribution $Q_0 \circ \Theta^{-n}$ on $\Omega \times E$. Suppose that the sequence $\{Q_n, n \geq 0\}$ is* tight. *Let Q be any subsequential weak limit of $\{Q_n, n \geq 0\}$. Then Q on $(\Omega \times E, \mathcal{F} \otimes \mathcal{E})$ with the shift Θ and with $X_0(\omega, x) = x$ and $\Phi_0(\omega, x) = \phi_0(\omega)$ is a weak solution of (20).*

Examples of non-monotone recursions that can be handled using this approach are discussed in [8]. The following uniqueness theorem is also useful, see [8]. Let $C_b(\Omega \times E)$ denote the space of bounded continuous functions on $\Omega \times E$.

Theorem 3.3 *Suppose that for every Θ-invariant probability distribution Q on $\Omega \times E$ having Ω marginal P, and all $f \in C_b(\Omega \times E)$*

$$\lim_{n \to \infty} \frac{1}{n} \sum_{j=0}^{n-1} f(\Theta^j(\omega, x)) = \mathcal{A}(f) \tag{23}$$

exists Q-a.s. and is a constant for Q-a.a. (ω, x). Then if there is a Θ-invariant probability distribution Q on $\Omega \times E$ having Ω marginal P, it is unique.
Conversely, if there is a unique Θ-invariant probability distribution Q on $\Omega \times E$ having Ω marginal P, then for all $f \in C_b(\Omega \times E)$ the limit in (23) exists Q-a.s.. and is constant for Q-a.a. (ω, x).

3.5 Kelly-type networks

We now turn to a discussion of Kelly-type networks, where the service time and routing variables are carried by the packets as they move around the network and queue in FCFS fashion at the nodes. The stability question for Kelly-type networks is very poorly understood. Indeed, already in special cases, examples are known where the rate conditions (11) do not guarantee stability. We now briefly sketch the genesis of these examples. Note that the first two examples outlined below are *not* Kelly-type networks.

The example of Figure 4 is described by Lu and Kumar [58], and attributed to Seidman; see also Kumar and Seidman [54]. There is a single server that can serve either of the buffers 1 and 4, and another server that can serve either of the buffers 2 and 3; each server works at rate 1. There are priority rules for which buffer the server can serve, as indicated in the figure : buffer 4 has priority over buffer 1, and buffer 2 has priority over buffer 3. Assume that there is an arrival at every integer time, and that its service requirements are 2/3 at each of the buffers 2 and 4, and that the service requirement is 0 at each of the buffers 1 and 3 : this means that the server just has to "kiss" the packet before allowing it to leave. Also assume that kissing at buffer 3 takes place just before kissing at buffer 1. Consider the initial condition (at time 0-) when there are M packets in buffer 1 and the other buffers are empty. All M inital packets are immediately kissed and go to buffer 2. Further, if we define the time τ by

$$\frac{2}{3}(\tau + M) = \tau \tag{24}$$

so that $\tau = 2M$, we see that, in view of the priority rules, at time $\tau-$ there will be 3M packets in buffer 3, and all the other buffers will be empty. These then get kissed at buffer 3, leaving 3M packets at buffer 4 with all the other buffers empty. Finally, in view of the priority rules, at time 4M- we are in a situation where there are 4M packets in buffer 1, with all the other buffers empty. The system has returned to a scaled version of the initial condition, with scaling bigger than 1.

What is remarkable about this example is that the total work that needs to be done per packet by each server is $\frac{2}{3} < 1$. Clearly rate conditions do not suffice to determine the stability region in this system with service priority rules.

The system above involves a fluid model, and instantaneous events. Motivated by a similar example in [54], Rybko and Stolyar [68], considered the model of Figure 5. There are two classes of packets. Assume that the respective arrival processes are independent Poisson processes of rate 1. The service time required by a packet in buffer ij is an exponential random variable of mean m_{ij}; all these service times are independent, and independent of the arrival processes. There is a server that works at rate 1 serving buffers 11 and 22, giving preemptive priority to buffer 22, and another server that works at rate 1 serving buffers 12 and 21, giving preemptive priority to buffer 12. The rate based condition for stability is then

$$\begin{aligned} m_{11} + m_{22} &< 1 \\ m_{12} + m_{21} &< 1 \,. \end{aligned}$$

In [68] the situation with $m_{12} = m_{22} = m_2 > \frac{1}{2}$ and $m_{11} = m_{21} = m_1 > 0$, and with $m_1 + m_2 < 1$ is considered. It is shown that the Markov process describing the system is *not* positive recurrent. Thus rate conditions do not suffice to determine stability in this example.

A sketch of the intuition behind this result is the following : Let $Q_{ij}(t)$ denote the queue size at buffer ij at time t, including the packet in service, if any. Then the network can described by a Markov process with state $(Q_{ij}(t), 1 \leq i, j \leq 2)$. To begin with, we may restrict attention to the subset of the state space such that $Q_{12}(t)Q_{22}(t) = 0$, since from any initial condition we reach such a state and then never leave this subset of states. This means that if $Q_{12}(t) > 0$, then $Q_{22}(t) = 0$, so that type 1 calls can be thought of as entering a tandem system of two rate 1 servers with infinite buffers in front of them. This picture is no longer valid when buffer 12 empties and a packet at buffer 11 is served and enters buffer 22 before a packet from buffer 11 is served and enters buffer 12. Now arriving type 1 packets build up till we have a situation where buffer 22 empties and a packet served at buffer 11 reaches buffer 12 before a packet served at buffer 21 enters buffer 22. Compressing the intervals where type 2 packets are being worked on gives the picture of type 1 packets entering a tandem system of two rate 1 servers but with an arrival process that includes large bursts of arrivals corresponding to the built up packets of type 1 during a cycle when type 2 packets are being worked on. It is now shown that these bursts destabilize the system. For details, see Section 6 of [68].

Both examples above are based on priority rules for the servers. Nevertheless, Bramson [13] had the insight that it is possible to mimic the phenomenon underlying these examples when the service discipline is FCFS, i.e. in Kelly-type networks. The invention of the examples of Bramson [13, 14], and fluid-model based examples by Seidman [70]. showing that the natural rate conditions (11) do not suffice to guarantee stability in Kelly-type networks is one of the most striking recent developments in the area of stochastic networks.

The example of [13] is illustrated in Figure 6. There is a single stream of customers entering the network according to a Poisson process of rate 1 and proceeding according to the route

$$\rightarrow 1 \rightarrow 2 \rightarrow 2 \rightarrow \ldots \rightarrow 2 \rightarrow 2 \rightarrow 1 \rightarrow \tag{25}$$

The service requirements at each visit along the route are exponentially distributed random variables, with mean c at the first visit to node 2 and the last (second) visit to node 1 and mean δ at the first visit to node 1 and all except the first visit to node 2. It is demonstrated in [13] that for c sufficiently close to 1, if the number of visits, J, to node 2 is sufficiently large (depending on c) and δ is sufficiently small (depending on c and J) the Markov process describing the system is transient. Note that this means that the natural rate conditions

$$\begin{aligned} c + \delta &< 1 \\ c + (J-1)\delta &< 1 \end{aligned}$$

do not suffice to guarantee stability. For details, see [13].

The second example of Bramson, presented in [14], demonstrates that for any given $\rho < 1$, however small, it is possible to have Kelly-type networks where the load factor at each node

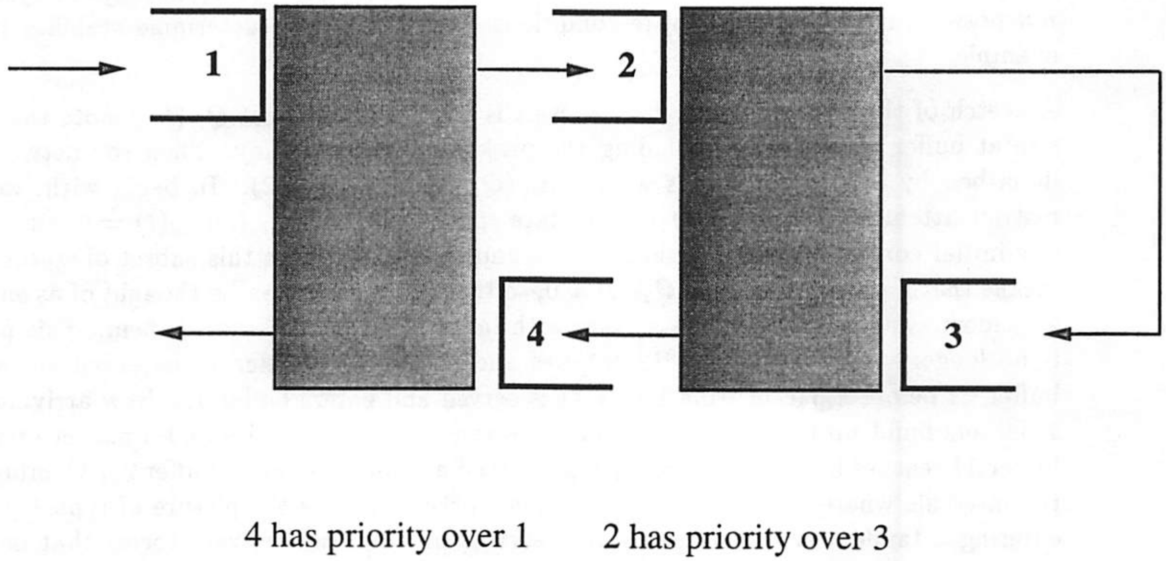

Figure 4: The Lu-Kumar Example

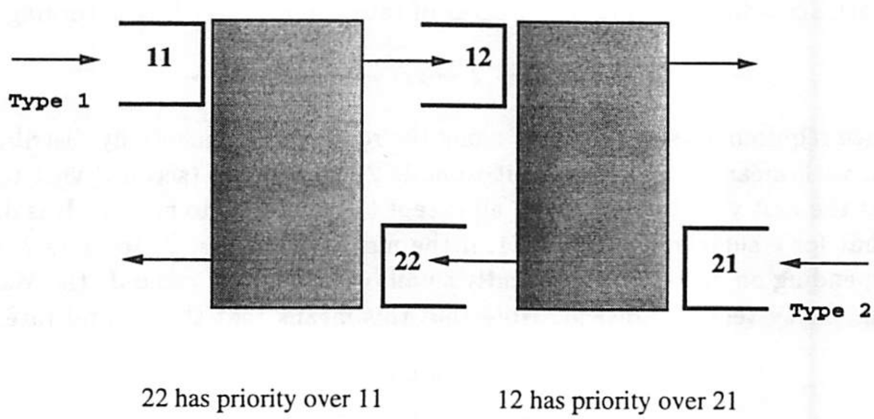

Figure 5: The Rybko-Stolyar Example

is less than ρ, but the network is nevertheless unstable. However, note that the smaller the prescribed ρ, the larger the number of nodes needed in order to construct an example of this sort.

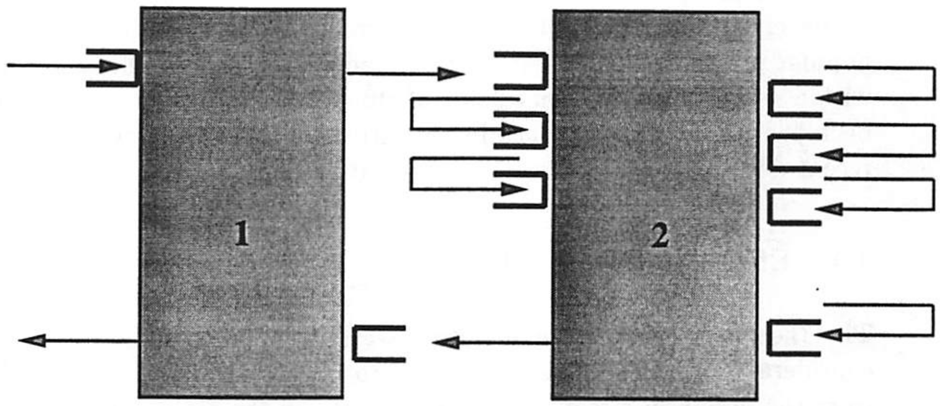

There is an enormous number of quick visits to node 2

Figure 6: Bramson's first example

Understanding the stability question for Kelly-type networks is one of the important and exciting challenges thrown up by the area of stochastic networks. Progress towards this problem in a Markovian framework is reported by Dai [25], Dai and Meyn [26], and Kumar and Meyn [53].

4 Integrated Broadband Networks

Network design engineers would like to have in place a network that can offer a wide variety of services, such as audio, video, data, etc. using a common protocol suite. Such a network is conceived of as the facilitator of the multimedia revolution that is booted about in the popular press. Progress in fibre optics and switching technology has made this goal appear within reach. From the viewpoint of stochastic analysis, a number of new questions have been brought into prominence by this drive. In this section we briefly introduce a couple of recent developments in research, aimed at addressing these questions.

4.1 Effective Bandwidths

The theory of effective bandwidths, see Hui [43] and Kelly [47], is currently a topic of considerable research, The rationale is to try to carry over some of the intuition available from the design of circuit-switched networks to the design of networks that handle bursty traffic. This is made possible by results in large deviations theory.

To convey the basic idea, we first recall the Gärtner-Ellis theorem, see [30, 36]. For more details, see Dembo and Zeitouni [27, Sec. 2.3]. Let $(X_n, n \geq 1)$ be a sequence of random variables and $S_n = \sum_{i=1}^{n} X_i$. Suppose that the limit

$$\Lambda(\theta) = \lim_{n \to \infty} \frac{1}{n} \log E[\exp(\theta S_n)] \tag{26}$$

exists (possibly infinite) for all $\theta \in \mathbf{R}$ and is lower semicontinuous. Define the effective domain of $\Lambda(\cdot)$ as $\{\theta \ : \ \Lambda(\theta) < \infty\}$. Suppose that $\theta = 0$ is in the interior of the effective domain, that $\Lambda(\cdot)$ is differentiable throughout the interior of its effective domain, and that the derivative approaches ∞ in absolute value for any sequence approaching a boundary point of the effective domain. Then $(\frac{S_n}{n}, n \geq 1)$ obeys a large deviations principle with convex good rate function $\Lambda^*(\cdot)$ given by the convex dual of Λ :

$$\Lambda^*(x) = \sup_{\theta}[\theta x - \Lambda(\theta)] \ , \tag{27}$$

i.e. for any closed set $F \subseteq \mathbf{R}$ we have

$$\limsup_{n \to \infty} \frac{1}{n} \log P(\frac{S_n}{n} \in F) \leq - \inf_{x \in F} \Lambda^*(x) \tag{28}$$

and for any open set $G \subseteq \mathbf{R}$ we have

$$\liminf_{n \to \infty} \frac{1}{n} \log P(\frac{S_n}{n} \in G) \geq - \inf_{x \in G} \Lambda^*(x) \ . \tag{29}$$

Here a rate function is said to be *good* if it has compact level sets.

Consider now the Lindley equation (13) with $X_n = \sigma_n - \tau_n$, which we reproduce here :

$$W_{n+1} = (W_n + X_n)^+ \ . \tag{30}$$

Then we may prove the following result, see [29].

Theorem 4.1 *Let $(X_n)_n$ be a stationary ergodic process with $E[X_n] < 0$, and satisfy the conditions of the Gärtner-Ellis theorem. Then there is a unique stationary solution to (30) which satisfies*

$$\Lambda(\theta) \leq 0 \equiv \lim_{B\to\infty} \frac{1}{B} \log P(W_n > B) \leq -\theta \ . \tag{31}$$

The probability that the stationary queue size exceeds some level is of particular interest in applications. Indeed, in practice buffer sizes are finite, and this probability can be taken as representative of the probability of buffer overflow. Since these probabilities are very small in well designed systems, an exponent of the form of the limit on the right hand side of (31) is of considerable interest. Theorem 4.1 gives a broad general connection between this exponent and the large deviations behaviour of the driving process $(X_n)_n$.

Consider next a collection of independent sources of different types $1, \ldots, J$. We think of time as divided into slots of identical length and assume that the work brought in by a traffic stream of type j in successive slots has the distribution of $(A^j_n)_n$, where $(A^j_n)_n$ satisfies the conditions of the Gärtner-Ellis theorem, with

$$\Lambda_j(\theta) = \lim_{n\to\infty} \frac{1}{n} \log E[\exp(\theta \sum_{i=1}^{n} A^j_n)] \ . \tag{32}$$

There are n_j traffic streams of type j, $1 \leq j \leq J$. These streams all share a buffer which is served by a work-conserving server that can serve at most c units of work during a slot. Assuming that the work brought in by a traffic stream arrives at the beginning of the slot, and letting W_n denote the total work in the buffer at the end of slot $n-1$, we see that $(W_n)_n$ obeys a Lindley equation of the form (30) with X_n being the total work brought in during slot n, less c.

For $\theta > 0$, let $\alpha_j(\theta) = \frac{\Lambda_j(\theta)}{\theta}$. The function $\alpha_j(\cdot)$ is called the *effective bandwidth* function of sources of type j. The reason behind this is the following theorem, obtained by Kesidis et al. [48], which follows from Theorem 4.1, see [29] for a proof.

Theorem 4.2 *For the buffer shared by n_j arrivals of each type j, $1 \leq j \leq J$, and served by a work conserving server that can serve at most c units of work during a slot, as above, we have*

$$\sum_j n_j \alpha_j(\theta) \leq c \equiv \lim_{B\to\infty} \frac{1}{B} \log P(W_n > B) \leq -\theta \ . \tag{33}$$

Thus the idea of effective bandwidth appears to convert problems of designing for small buffer overflow probability into problems similar to fitting calls on a link in a circuit switched network. Much of the current work in this area is aimed at trying to carry over this analogy to networks of queues. There has also been considerable interest in computing formulas for the effective bandwidths of different source types. For a better perspective on this rapidly evolving area see [19, 31, 64].

4.2 Virtual backlog

Before a traffic stream is admitted to the network it is regulated to prevent its potential burstiness from adversely affecting the ability of the network to handle other traffic streams sharing the network. This process is called *flow control.* One of the features of broadband networks is the difficulty of effectively using feedback for flow control. This is because a large number of packets would have already entered the network before feedback has a chance to come into play. For this reason the flow control schemes currently being proposed are largely open loop in nature. Of these, variants of the *leaky bucket* flow control scheme, see [75], are by far the most popular.

Broadband networks are likely to use the ATM (Asynchronous Transfer Mode) protocol suite, see [28, 40]. In an ATM network, traffic is broken up into packets of fixed length, called *cells.* The basic idea of the leaky bucket scheme is to regulate the admission of cells into the network by means of a stream of tokens generated at a constant rate with constant interarrival times. The tokens collect in a token buffer, and the cells collect in a cell buffer. A cell is released only when there is a token available, in which case it consumes one token. Tokens arriving when the token buffer is full are lost, as are cells arriving when the cell buffer is full. This scheme has been found to be very effective in reducing the burstiness of offered traffic in practice, and this burstiness reducing property has also been analytically justified, see [7, 51]. For instance, Anantharam and Konstantopoulos [7] prove the following result : Consider a stationary point process as bringing one unit of work with each point. Define a stationary point process to be less bursty than another such process if the steady state queue length in any single server queue working at rate strictly bigger than the total rate of arrival of work in either process is stochastically smaller than that for the other process. Then for *any* stationary arrival process of offered traffic into a leaky bucket flow controller, if the token arrival rate is at least as big as the rate of offered traffic, the burstiness of the stationary departure process is monotonically increasing in the size of the token buffer. Here we assume that the cell buffer has infinite capacity. Note that the degenerate case of infinite token buffer corresponds to not regulating the arrival process at all, so this result shows that the departure process from the leaky bucket (the traffic admitted to the network) is less bursty than the process of arrivals into the leaky bucket (the offered traffic from the source).

We observe that the leaky bucket scheme regulates offered traffic so that it satisfies *burstiness constraints* of the form that the amount of traffic over a time interval is bounded by an affine function of the length of the interval. Indeed, if the size of the token buffer is C, and the rate of token generation is ρ, the total number of cells that can be released by the bucket over a time interval $[a, b]$ can be no more than $C + 1 + \rho(b - a)$. The class of traffic flows that satisfy such burstiness constraints is therefore an interesting one to consider. There has been a considerable amount of work on such models of traffic flow, beginning with Cruz [22, 23]; see also Chang [19], and Parekh and Gallager [65, 66].

A key observation is that it is possible to characterize the past of a burstiness constrained flow in terms of a simple recursively updatable statistic which we call the *virtual backlog.* To make this precise let us assume, for convenience, that time is discrete, although similar ideas carry over to continuous time. We will use the term *message flow* to denote a sequence of nonnegative real numbers $(a_n, n \geq 0)$. We model traffic on the links of a network by a message flow. A message flow $(a_n, n \geq 0)$ is said to be (σ, ρ) constrained, if for all $0 \leq n_0 \leq n_1 < \infty$, we have

$$\sum_{n_0}^{n_1} a_k \leq \sigma + \rho(n_1 - n_0 + 1) \,. \tag{34}$$

Let $(a_n, n \geq 0)$ be a (σ, ρ) constrained flow. We say that it has initial *virtual backlog* σ_0, if in addition to (34) the flow obeys the constraints

$$\sum_{0}^{n} a_k \leq \sigma_0 + \rho(n+1) \tag{35}$$

for all $n \geq 0$. Then we may easily prove the following result, see [4] for instance.

Lemma 1 *Let $(a_n, n \geq 0)$ be a (σ, ρ) constrained flow with initial virtual backlog σ_0. Suppose that a_0 is revealed. Then the information gained about $(a_n, n \geq 1)$ is completely summarized by the statement that $(a_n, n \geq 1)$ is a (σ, ρ) constrained flow with initial virtual backlog σ_1, where*

$$\sigma_1 = \min(\sigma_0 + \rho - a_0, \sigma) \tag{36}$$

This observation allows one to adopt a system theoretic viewpoint to the design of resource allocation and control strategies in networks handling burstiness constrained flows. Indeed, the virtual backlog is a *recursively updatable state* that *completely summarizes* the past of a burstiness constrained flow. We have been able to use this intuition to approach the problem of designing optimal open loop flow control strategies for traffic in broadband networks from a prescriptive point of view. For instance, Konstantopoulos and Anantharam [49] pose and solve the problem of regulating an arbitrary traffic stream to create a (σ, ρ) constrained traffic stream, while subjecting it to minimal delay. The permissible schemes are allowed in principle to use the *entire* past history of the offered traffic. The possibility of additional constraints such as a limit on the amount of traffic that can be buffered or a limit on the acceptable delay of traffic is also considered. The optimal control schemes that we find are very simple in nature (in fact they are greedy schemes), and are based on the virtual backlog of the admitted traffic stream. They can therefore be implemented with very little intelligence at the flow controller. For the problem with a constraint on the amount of traffic that can be buffered the optimal solution in fact turns out be a leaky bucket flow controller.

We now describe an abstract point of view for the control problems faced by network elements inside a network handling burstiness constrained traffic, relying on the idea of virtual backlog, see [4]. This idea was first discussed in [3]. We visualize the network element as being fed by a number of burstiness constrained flows and as implementing actions based on past information. Since we are working in discrete time throughout, we will assume that the *state* of the network element at time n, $n = 0, 1, \ldots$ is given by an element $\xi_n \in \Xi$. At time n the network element is to choose a *control action* $u_n \in U$. At a switch, for instance, U might represent a choice of matching between input ports and output ports. The evolution of the state of the network element occurs in response to the incoming flows at the current time and the choice of control action, resulting in the abstract evolution equation

$$\xi_{n+1} = f(\xi_n, u_n, \underline{a}_n) \tag{37}$$

Note that when there are K driving message flows with the the i th flow being (σ^i, ρ^i) constrained, $1 \leq i \leq K$, the domain of f is $\Xi \times U \times \prod_{i=1}^{K} [0, \sigma^i + \rho^i]$. A *randomized adapted control strategy* at the network element is a choice of a probability distribution on U at each time n as a function of $\xi_{[0,n]}$, $u_{[0,n-1]}$, $\underline{\sigma}_0$, and $\underline{a}_{[0,n-1]}$. Here $\underline{\sigma}_0$ is the vector of initial virtual backlogs of the flows.

The problem of designing a good control strategy can be posed via the theory of zero-sum stochastic games. For instance, we may postulate a function $c(\xi, u, \underline{a})$ representing the cost associated to taking the action u when the current element state is ξ and the current message flows are given by $\underline{a}$. We think of this as a cost *paid by the controller to the sources.* Let $0 < \beta < 1$ be a discount factor. We may then formulate the problem of the controller as one of minimizing the total infinite horizon discounted cost

$$\sum_{n=0}^{\infty} \beta^n c(\xi_n, u_n, \underline{a}_n) \tag{38}$$

The minimization is to be done over all possible randomized adapted control strategies of the controller and over all possible randomized adapted control strategies of the sources. Since the source action space is convex, this is equivalent to a worst case formulation where the controller attempts to minimize the overall discounted cost over all possible burstiness constrained source sequences.

Using techniques that are fairly standard in the theory of stochastic games, one can prove the following theorem. For the basic ideas behind the Shapley recursion, see Shapley [72].

Theorem 1 *Suppose Ξ is a complete separable metric space and U is compact. Suppose f is continuous. Suppose there is a continuous nonnegative function g defined on Ξ having compact level sets, a polynomial $p(\cdot)$, and a constant $K < \infty$ such that, for all $\xi \in \Xi$, $u \in U$, and $\underline{a} \in \prod_{i=1}^{K} [0, \sigma^i + \rho^i]$*

(1) $|g(f(\xi, u, \underline{a})) - g(\xi)| \leq K$.

(2) $c(\xi, u, \underline{a}) \leq p(g(\xi))$.

Then the stochastic game above admits a continuous value function that is the unique fixed point of a Shapley recursion. Further, both the controller and the sources have optimal stationary randomized adapted strategies which depend only on the state of the game. These are respectively given by the outer extremizers in the min-max and the max-min forms of the Shapley recursion.

By iterating the Shapley recursion, optimal control strategies for the controller within the context of this formulation can be identified. Further, they can *in principle be implemented in real time.* Indeed, the entire history of a burstiness constrained flow can be kept track of by the *recursively updatable* virtual backlog, which is a simple finite dimensional statistic. This therefore appears to be a promising approach to handle resource allocation problems in broadband networks. Analytical results about the qualitative nature of optimal control strategies in specific problems would also be welcome. In practice, of course, one would expect that the basic burstiness parameters should be updated on a slower time scale.

5 Concluding remarks

We have given a rather selective sketch of some of the recent research in the area of stochastic networks, organizing them along the lines of distinct applications contexts. Large areas of highly interesting work, both from the viewpoint of mathematics and from the viewpoint of applications, have been left entirely untouched; this includes work on diffusion approximations, see [21, 41, 24, 42]; on the use of self-similar traffic models, which appear to match actual samples of local area network traffic much better than Markovian models, see Leland et al. [55]; on the analysis of different scheduling policies for re-entrant lines, see Kumar [52]; on the stochastic perfomance analysis of interconnection networks that form the switching fabric of high speed networks, see [67]; among many others. The reader whose appetite is whetted by this survey, and who follows the track of some of the references, will no doubt quickly find a wealth of interesting and useful problems to occupy his or her attention.

References

[1] Akinpelu, J. M. (1984). The Overload Performance of Engineered Networks with Non-hierarchical and Hierarchical Routing. *AT & T Bell Laboratories Technical Journal*, Vol. 63, No. 7, pp. 1261 -1281.

[2] Anantharam, V. (1991). A Mean Field Limit for a Lattice Caricature of Dynamic Routing in Circuit Switched Networks. *Annals of Applied Probability*, Vol. 1, No. 4, pp. 481 -503.

[3] Anantharam, V. (1992). Bandwidth allocation to burstiness constrained flows. Talk delivered at the *IEEE Information Theory Workshop*, Salvador, Brazil, June 1992.

[4] Anantharam, V. (1993). An approach to the design of high speed networks for bursty traffic. *Proceedings of the International Teletraffic Congress sponsored seminar on Teletraffic Methods for Current and Future Telecom Networks*, Indian Institute of Science, Bangalore, India, Nov. 15 -19, 1993, pp. 163 -168.

[5] Anantharam, V. (1994). Large deviations for a Mean Field Model of Interacting Particles. *In preparation.*

[6] Anantharam, V., and Benchekroun, M. (1993). A Technique for Computing Sojourn Times in Large Networks of Interacting Queues. *Probability in the Engineering and Information Sciences*, Vol. 7, pp. 441 -464.

[7] Anantharam, V., and Konstantopoulos, T. (1994). Burst reduction properties of the leaky bucket flow control scheme in ATM networks. *IEEE Transactions on Communications*, Vol. 42, No. 12, pp. 3085 -3089.

[8] Anantharam, V., and Konstantopoulos, T. (1994). Stationary Solutions of Stochastic Recursions describing Discrete Event Systems : General results. *Technical Report*, No. SCC-94-03, ECE Department, University of Texas, Austin.

[9] Baccelli, F., and Brémaud, P. (1994). *Elements of Queueing Theory : Palm-Martingale Calculus and Stochastic Recurrences.* Springer-Verlag, Berlin.

[10] Baccelli, F., and Foss, S. (1994). Ergodicity of Jackson-type queueing networks. *Queueing Systems : Theory and Applications*, Vol. 17, pp. 5 -72.

[11] Baccelli, F., and Foss, S. (1994). On the Saturation rule for the Stability of Queues. *Preprint. To appear in Journal of Applied Probability.*

[12] Borovkov, A. A. (1986). Limit Theorems for Queueing Networks. *Theory of Probability and its Applications*, Vol. 31, pp. 413 -427.

[13] Bramson, M. (1994). Instability of FIFO queueing networks. *Annals of Applied Probability*, Vol. 4, No. 2, pp. 414 -431. Correction published in Vol. 4, No.3, pg. 952.

[14] Bramson, M. (1994). Instability of FIFO queueing networks with quick service times. *Annals of Applied Probability*, Vol. 4, No. 3, pp. 693 -718.

[15] Brandt, A., Franken, P., and Lisek, B. (1990). *Stationary stochastic models.* John Wiley and Sons, New York.

[16] Brémaud, P., Kannurpatti, R., and Mazumdar, R. (1992). Event and Time Averages : A review. *Annals of Applied Probability*, Vol. 24, pp. 377 -411.

[17] Burke, P. J. (1956). The output of a queueing system. *Operations Research*, Vol. 4, pp. 699 -704.

[18] Burman, D. Y., Lehoczy, J. P., and Lim, Y. (1984). Insensitivity of Blocking Probabilities in a Circuit-switching Network. *Journal of Applied Probability*, Vol. 21, pp. 853 -859.

[19] Chang, C.-S. (1994). Stability, Queue Length, and Delay of Deterministic and Stochastic Queueing Networks. *IEEE Transactions on Automatic Control*, Vol. 39, No. 5, pp. 913 -931.

[20] Chang, C.-S., Thomas, J. A., and Kiang, S. H. (1994). On the stability of open networks : a unified approach by stochastic dominance. *Queueing Systems : Theory and Applications*, Vol. 15, pp. 239 -260.

[21] Chung, K.L. and Williams, R.J. (1983). *Introduction to Stochastic Integration.* Birkhäuser, Boston.

[22] Cruz, R.L. (1991). A calculus for network delay, part I: Network elements in isolation. *IEEE Transactions on Information Theory*, Vol. 37, 114-131.

[23] Cruz, R.L. (1991). A calculus for network delay, part II: Network analysis. *IEEE Transactions on Information Theory*, Vol. 37, 132 -141.

[24] Dai, J. G., and Wang, Y. (1993). Nonexistence of Brownian models for certain multiclass queueing networks. *Queueing Systems : Theory and Applications*, Vol. 13, pp. 41 -46.

[25] Dai, J. G. (1994). On Positive Harris Recurrence for Multiclass Queueing Networks : A Unified Approach via Fluid Limit Models. *Preprint.*

[26] Dai, J., and Meyn, S. (1994). Stability and Convergence of Moments for Multiclass Queueing Networks via Fluid Limit Models. *Preprint.*

[27] Dembo, A., and Zeitouni, O. (1993). *Large Deviations : Techniques and Applications.* Jones and Bartlett Publishers, London.

[28] de Prycker, M. (1993). *Asynchronous Transfer Mode : Solution for Broadband ISDN.* Ellis Horwood, New York.

[29] de Veciana, G., and Walrand, J. (1993). Effective bandwidths : call admission, traffic policing, and filtering for ATM networks. *Preprint. To appear in the IEEE/ACM Transactions on Networking.*

[30] Ellis, R. S., (1984). Large deviations for a general class of random vectors. *Annals of Probability*, Vol. 12, pp. 1 -12.

[31] Elwalid, A. I., and Mitra, D. (1993). Effective bandwidth of general Markovian traffic sources and admission control of high-speed networks. *IEEE/ACM Transactions on Networking*, Vol. 1, pp. 329 -343.

[32] Erlang, A. K. (1917). Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *Post Office Electrical Engineers Journal*, Vol. 10, pp. 189 -197.

[33] Ethier, S. N., and Kurtz, T. G. (1986). *Markov Processes: Characterization and Convergence* John Wiley and Sons.

[34] Foss, S. G. (1989). Some properties of open queueing networks. *Problems of Information Transmission*, Vol. 25, pp. 241 -246.

[35] Foss, S. G. (1991). Ergodicity of Queuing Networks. *Siberian Mathematical Journal*, Vol. 31, pp. 184 -203.

[36] Gärtner, J. (1977). On large deviations from the invariant measure. *Theory of Probability and its Applications*, Vol. 22, pp. 24 -39.

[37] Gibbens, R. J., Hunt, P. J., and Kelly F. P. (1990). Bistability in Communications Networks. In *Disorder in Physical Systems*, Grimmett, G. R., and Welsh, D. J. A., eds., Oxford University Press, pp. 113 -128.

[38] Graham, C., and Méléard, S. (1993). Propagation of Chaos for a fully connected loss network with alternate routing. *Stochastic Processes and their Applications*, Vol. 44, pp. 159 -180.

[39] Graham, C., and Méléard, S. (1993). Fluctuations for a fully connected loss network with alternate routing. *Preprint.*

[40] Händel, R., Huber, M. N., and Schröder, S. (1994). *ATM Networks : Concepts, Protocols, Applications.* Addison-Wesley, Reading, Massachusetts.

[41] Harrison, J.M. (1985). *Brownian Motion and Stochastic Flow Systems.* Wiley, New York.

[42] Harrison, J.M., and Nguyen, V. (1993). Brownian models of multiclass queueing networks : Current status and open problems. *Queuing Systems : Theory and Applications*, Vol. 13, pp. 5 -40.

[43] Hui, J. (1988). Resource allocation for broadband networks. *IEEE Transactions on Selected Areas in Communications*, Vol. 6, pp. 1598 -1608.

[44] Jackson, J. R. (1963). Jobshop-like queueing systems. *Management Science*, Vol. 10, pp. 131 -142.

[45] Kelly, F. (1986). Blocking probabilities in large circuit-switched networks. *Advances in Applied Probability*, Vol. 18, pp. 473 -505.

[46] Kelly, F. (1991). Loss Networks. Special Invited paper in the *Annals of Applied Probability*, Vol. 1, No. 3, pp. 319 -378.

[47] Kelly, F. (1991). Effective Bandwidths at multi-class queues. *Queueing Systems : Theory and Applications.* Vol. 9, Nos. 1-2, pp. 5 -16.

[48] Kesidis, G., Walrand, J., and Chang, C.-S. (1993). Effective Bandwidths for Multiclass Markov Fluids and Other ATM Sources. *IEEE/ACM Transactions on Networking*, Vol. 1, No.4, pp. 424 -427.

[49] Konstantopoulos, T., and Anantharam, V. (1993). Optimal Flow Control Schemes that Regulate the Burstiness of Traffic. *SCC Technical Report SCC-93-15*, ECE Department, University of Texas, Austin. (submitted to *IEEE Transactions on Communications*).

[50] Krengel, U. (1985). *Ergodic Theorems.* de Gruyter Studies in Mathematics, No. 6, Walter de Gruyter, Berlin.

[51] Kuang, L. (1994) On the variance reduction property of the buffered leaky bucket. *IEEE Transactions on Communications*, Vol. 42, No. 9, pp. 2670-2671.

[52] Kumar, P. R. (1993). Re-entrant lines. Queueing Systems : Theory and Applications, Vol. 13, Nos. 1 -3, pp. 87 -110.

[53] Kumar, P. R., and Meyn, S. (1993). Duality and Linear Programs for Stability and Performance Analysis of Queueing networks and Scheduling Policies. *Technical Report*, University of Illinois, Urbana. To appear in *IEEE Transactions on Automatic Control.*

[54] Kumar, P. R., and Seidman, T. (1990). Dynamic Instabilities and Stabilization Methods in Distributed Real-Time Scheduling of Manufacturing Systems. *IEEE Transactions on Automatic Control*, Vol. 35, pp. 289 -298.

[55] Leland, W. E., Taqqu, M. S., Willinger, W., and Wilson, D. V. (1994). On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, Vol. 2, No. 1, pp. 1 -15.

[56] Liggett, T. (1985). *Interacting Particle Systems.* Grund. der Math. Wiss., Vol. 276, Springer-Verlag, Berlin.

[57] Loynes, R. M. (1962). The stability of queues with non independent inter-arrival and service times. *Proceedings of the Cambridge Philosophical Society*, Vol. 58, pp. 497 -520.

[58] Lu, S. H., and Kumar, P. R. (1991). Distributed scheduling based on due-dates and buffer priorities. *IEEE Transactions on Automatic Control*, Vol. 36, pp. 1406 -1416.

[59] Marbukh, V. V. (1985). Asymptotic Investigation of a Complete Communications Network with a large number of points and bypass routes. *Problems of Information Transmission*, pp. 212 -216.

[60] Marbukh, V. V. (1985). Fully Connected Message-Switching Network with a Large number of Nodes, Bypass routes and a limited number of Waiting places at the nodes. *Problems of Information Transmission*, pp. 154 -161.

[61] Meyn, S. P., and Down, D. (1994). Stability of generalized Jackson networks. *Annals of Applied Probability*, Vol. 4, No. 1, pp. 124 -148.

[62] Meyn, S. P., and Tweedie, R. L. (1993). *Markov chains and Stochastic Stability.* Springer-Verlag, London.

[63] Nelson, R. D. (1993). The Mathematics of Product-Form Queuing Networks. *Computing Surveys*, Vol. 25, No. 3, pp. 339 -369.

[64] O'Connell, N. (1994). Large deviations in Queueing Networks. *Preprint.*

[65] Parekh, A. K., and Gallager, R. G. (1993). A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks : The Single Node Case. *IEEE/ACM Transactions on Networking*, Vol. 1, pp. 344 -357.

[66] Parekh, A. K., and Gallager, R. G. (1994). A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks : The Multiple Node Case. *IEEE/ACM Transactions on Networking*, Vol. 2, No. 2, pp. 137 -150.

[67] Robertazzi, T. G., (Editor) (1993). *Performance Evaluation of High Speed Switching Fabrics and Networks : ATM, Broadband ISDN, and MAN Technology.* IEEE Press, New York.

[68] Rybko, A., and Stolyar, A. L. (1992). Ergodicity of stochastic processes that describe the functioning of of open queueing networks. *Problems of Information Transmission*, Vol. 28, pp. 3 -26.

[69] Schwartz, M. (1987). *Telecommunication Networks : Protocols, Modeling and Analysis.* Addison-Wesley, Reading, Massachusetts.

[70] Seidman, T. (1994). First-come first-served is unstable. *Preprint.*

[71] Shantikumar, J. G., and Yao, D. (1989). Stochastic monotonicity in general queueing networks. *Journal of Applied Probability*, Vol. 26, pp. 413 -417.

[72] Shapley, L. (1953). Stochastic Games. *Proceedings of the National Academy of Sciences*, Vol. 39, pp. 1095 -1100.

[73] Sigman, K. (1990). The stability of open queueing networks. *Stochastic Processes and Their Applications*, Vol. 35, pp. 11 -25.

[74] Sznitman, A. (1989). Propagation of chaos. In *Ecole d'Eté de Probabilités de Saint-Flour*, Lecture Notes in Mathematics, Vol. 1464, Springer-Verlag, Berlin.

[75] Turner, J. (1986). New Directions in Communications (or Which Way in the Information Age). *IEEE Communications Magazine*, Vol 24, pp. 8 -15.

[76] Uchiyama, K. (1985). Fluctuation in Population Dynamics. In *Lecture Notes in Biomathematics*, Vol. 70, pp. 222-229.

[77] Walrand, J. (1983). A Probabilistic Look at Networks of Quasireversible Queues. *IEEE Transactions on Information Theory*, Vol. IT-29, pp. 825 -831.

[78] Walrand, J. (1988). *An Introduction to Queueing Networks.* Prentice Hall, Englewood Cliffs, N.J.

[79] Wentzell, A.D., and Freidlin, M. I. (1984). *Random perturbations of dynamical systems.* Grundlehren der mathematischen Wissenschaften, Vol. 260, Springer-Verlag, Berlin.

[80] Wolff, R. W. (1982). Poisson arrivals see time averages. *Operations Research*, Vol. 30, pp. 223 -231.