STABILITY OF COMMUNICATION NETWORKS

Aman Kumar Maulloo

UNIVERSITY OF MAURITIUS

A dissertation submitted for the degree of Doctor of Philosophy at the University of Mauritius

A major part of the research was conducted at the Statistical Laboratory, University of Cambridge, UK.

April 15, 2000

Contents

List of Figures

Acknowledgement

I would like to acknowledge all the help and encouragement I have received in the course of reading for my PhD. I am especially grateful to my external supervisor, Professor F.P. Kelly, F.R.S., whose work in the dynamic field of communication systems modelling inspired me and without whose persistent encouragement and guidance throughout this networking experience, the writing up of this thesis would not have been so challenging.

For his continuous support and advice, I am indebted to my local supervisor, Professor M.K. Jain, D.Litt., D.Sc. The support of the University of Mauritius and its Staff Development Programme is also acknowledged: without their financial assistance, I would not have been able to read for my PhD in Cambridge. Thanks are also due to many of my colleagues of the University of Mauritius, especially those of the Department of Mathematics for their support as and when required.

A special mention for all members of the Statistical Laboratory, University of Cambridge, for providing a sound and healthy research environment during my visits to this prestigious institution.

I am also grateful to Mr and Mrs Tanesh Bhugobaun and family who have provided me with the necessary support to make my stay in Cambridge smooth running.

Last but not least Mr and Mrs K. Heywood and my inmates at Link House, for making my final year very enjoyable indeed.

I dedicate this dissertation to my wife, who has been the main driving force behind this research work.

Declaration

I certify that this thesis does not incorporate without acknowledgement any material previously submitted for a degree or diploma in any University. The work in the second chapter was conducted in collaboration with Professor F. P. Kelly, F. R. S., and D. K. H. Tan of the University of Cambridge. The joint work has been published in the *Journal of the Operational Research Society*, **Volume 49**, pages 237–252, 1998. The remaining content of the dissertation is the result of my own effort; to the best of my knowledge and belief, it does not contain any material previously published or written by another person where due reference is not made.

Aman Kumar Maulloo
Department of Mathematics
University of Mauritius

Prof. M.K. Jain, D.Litt., D.Sc.
Department of Mathematics
University of Mauritius

Summary

This dissertation is motivated by the significant growth in size and scope of large-scale broadband communication networks such as the Internet. We consider the problem of congestion in these networks and study how competing streams of elastic traffic may have a fair share of the available bandwidth while keeping the system stable.

We first present an optimisation framework for the control of the overall system which leads to a decomposition into the user and network subproblems respectively. We then analyse the stability and fairness of a primal-dual pair of rate control algorithms, which may be interpreted in terms of either congestion indication feedback signals or explicit rates based on shadow prices. These schemes are shown to be stable about a system optimum characterised by a proportional fairness criterion. The models are further extended to include two important aspects of communication networks, namely stochastic effects and time-lags.

The rate control algorithms are designed for large-scale networks which may not have definite structures due to the fast evolving technological advances in that area. This motivates us to explore the behaviour of our schemes for networks with random topologies. We derive some simple results for the eigenvalues of the linearised system in terms of the model parameters. These explicit forms allow us to control the rate of convergence to steady state as well as the fluctuations by proper choice of the parameters, either by the end-users, or by the network resources.

A new approach to flow control is to allow the end-users to choose the parameters in order to make efficient use of the network. With this reasoning, we study an algorithm related to TCP/IP for a network with a single bottleneck resource. The concept of packet marking is used to indicate congestion for a set of parameters which can be scaled to suit the realistic models. We show how the choice of the contract parameters by the end-users, in particular, the gain parameters, affect stability at and convergence times to equilibrium. The results are illustrated by the simulation of different scenarios.

Chapter 1

Introduction

The design and control of modern communication networks raise several issues well suited to study using techniques of operational research such as optimisation, network programming and stochastic modelling. In this dissertation, we illustrate this theme, through the presentation and analysis of a mathematical model that arises in connection with the development and deployment of large-scale broadband networks.

In future communication networks, there are expected to be applications that are able to modify their data transfer rates according to the available bandwidth within the network. Traffic from such applications is termed *elastic* [89]; a typical current example is the Transfer Control Protocol (TCP) traffic over the Internet [34, 57], and future examples may include the controlled-load service of the Internet Engineering Task Force (IETF) and the Available Bit Rate (ABR) transfer capability of the Asynchronous Transfer Mode (ATM) networks [56]. In recent years, the Internet has grown significantly in size and scope, and as a result new protocols and algorithms are being developed to meet changing operational requirements in the Internet [10]. Although the current TCP congestion control has been relatively successful, its ability is exceedingly stretched by the rapid growth of the Internet and the proliferation of both real-time and multicast services. Over the past several years, considerable effort has been directed at improving the existing techniques of congestion control in the Internet and at introducing new approaches to accommodate the requirements of new services and applications [48]. Floyd [35] and Jain [58] provide a good tutorial on congestion control principles.

The key issue we address in this dissertation concerns how the available bandwidth within a network should be shared among competing streams of elastic traffic; in particular, we present a tractable mathematical model and use it to analyse the stability and fairness of two classes of rate control algorithms. Two important aspects of decentralised systems, namely stochastic perturbations and time-lags, will also be considered.

Traditionally, stability has been considered an engineering issue, requiring an analysis of randomness and feedback operating on fast time-scales, while fairness and pricing have been considered economic issues, involving static comparisons of utility. In future networks, the intelligence embedded in end-systems, acting on behalf of human users, is likely to lessen the distinction between engineering and economic issues and increase the importance of an interdisciplinary view (for further discussion, see [64]).

Economists have suggested using pricing as a means of avoiding congestion in the Internet. In this dissertation, we have assumed that pricing may be used to control demand and hence congestion in the network, and do not intend to discuss pricing in detail. In [36], Floyd and Fall discussed the possibility that unresponsive flows could contribute to congestion collapse in the Internet, since it is no longer possible to rely on all end-nodes to use end-to-end congestion control, or on developers to incorporate end-to-end congestion control in their applications. Mackie-Mason and Varian [76] showed that *flat-pricing* could lead to either higher congestion or lower capacity at equilibrium, both of which are undesirable.

This is already evident in the case of the Internet, and Odlyzko [84] has observed the evidence of repressed demand from available traffic statistics. (For more discussion on pricing, the reader is referred to the collection of papers in [79].)

Several pricing mechanisms have been proposed to control congestion in broadband networks. Odlyzko [85] proposed the Paris Metro Pricing (PMP); a simple mechanism, in which the network is partitioned into two sub-networks. Both sub-networks offer no special distinctive guarantees on the connection but they charge different prices for usage. Gibbens et al. [46] have shown that this scheme works well in a monopoly but may fail when there is competition in the market.

In a shared environment, the throughput for a source depends upon the demands by the other sources [58]. If there exists more than one connection sharing a *bottleneck* link, it is reasonable to ensure that the various user connections are being treated in an equitable way relative to each other. However, given the limited amount of information currently available at the connection end-points, the proper share of the bandwidth is unlikely to take place without some help from the routers or resources [17, 24]. This gives rise to the issue of *fairness* as a measure for judging the performance of a network [88]. According to Chiu and Jain [22], it is one of the four criteria of choosing a control algorithm; the others being efficiency, distributedness, and convergence. The issue of fairness is essential since there exist situations where a given scheme might maximise network throughput, for example, while denying access to some users. Fairness has been a key issue in many recent works such as [13, 58, 77].

The most commonly discussed fairness criterion is that of *max-min* fairness: loosely, a set of rates is max-min fair if no rate may be increased without simultaneously decreasing another rate which is already smaller. Mazumdar et al. [78] have pointed out that from a game-theoretic standpoint such an allocation is not special, and have advocated instead the Nash bargaining solution, from cooperative game theory, as capturing natural assumptions as to what constitutes fairness. It is a classical sharing principle in the domain of data networks as discussed by Bertsekas and Gallager [13] and is adapted by the ATM forum for congestion management of ABR service [21]. The reader is referred to [13, 19, 20, 59, 77, 88] for some examples on max-min fairness.

Max-min fairness is achieved by explicit rate calculated algorithms such as those studied in ABR context of ATM, as explained in Hernandez and Valencia [52]. In [19], Charny proposed an algorithm in which rate allocation is performed by the switches and calculated rates explicitly communicated to the sources in packet headers. This algorithm was shown by Charny et al. in [20] to converge to an exact max-min fair rate allocation in a finite number of iterations. They also verified that a rate allocation which is max-min fair is also *efficient* in the sense that given the fair allocation, no more data can be pushed through the network, since each route crosses at least one fully saturated link. Assuming that packets are infinitely small and flows are deterministic, it can be seen that max-min fairness implies maximum efficiency in the sense that the bottleneck resource is utilised up to its capacity and no queues build up. In the past decade, a number of algorithms have been developed to achieve max-min fair rates, see for example [13, 19, 54, 57, 83].

Although max-min fairness shows some significant advantages, there is no clear economic reason why max-min sharing should be preferred over some other bandwidth allocation [77]. This fairness criterion is imperfectly achieved by most network flow control protocols such as the additive increase and multiplicative decrease algorithm for the TCP [57]. Moreover, a max-min fair vector needs global information and most of these algorithms require exchange of information between the network and the hosts. More recently, Mo et al. [82] showed that while TCP Vegas [17] does achieve fairness, TCP Reno [30] fails to achieve fairness since TCP is not a synchronised rate based control scheme. The appropriateness of max-min fairness as a bandwidth sharing objective has recently been questioned by Kelly [66] who has introduced the alternative notion of *proportional fairness*.

The need for networks to operate in a public (and therefore potentially non-co-operative) environment has stimulated work on charging schemes for broadband networks: see [63] for a scheme based on time and volume measurements for non-elastic traffic. Mackie-Mason and Varian [76] provide a description of a *smart market*, which is an auction mechanism based on a per-packet charge when the network is congested. The reader is also referred to [44] and [45], where the authors discuss a simple pricing mechanism providing end-nodes with the necessary information and correct incentive to use a network efficiently. In [66], Kelly proposed a model in which a user chooses the charge per unit time that he is willing to pay; thereafter the user's rate is determined by the network according to a proportional fairness criterion applied to the rate per unit charge.

A vector $x = (x_r, r \in R)$ is proportionally fair if it maximises $\sum_{r \in R} \log x_r$ and it is feasible, that is $x \geq 0$ and $Ax \leq C$. It is unique in the case of finitely many links and routes. The objective may be interpreted as maximising the overall utility of rate allocations assuming each route has a logarithmic utility function (the law of diminishing returns). It may be characterised as follows: the aggregate of proportional rate charges with respect to the optimum of any other feasible vector $x^* = (x^*_r, r \in R)$ is negative

$$\sum_{r \in R} \frac{x^*_r - x_r}{x_r} \leq 0. \tag{1.1}$$

In a linear network, Roberts and Massoulié [88] show that max-min fairness is preferable to proportional fairness under random traffic while the converse is true under the assumption of a static configuration of persistent flows. The additive increase/multiplicative decrease congestion avoidance principle [22] as implemented for instance in TCP [57], tends to favour proportional rather than max-min fairness [54]. The interested reader is referred to [77, 88] for more details on this concept of fairness.

It was shown in [66] that a system optimum is achieved when users' choices of charges and the network's choice of allocated rates are in equilibrium. There remained the question of how the proportional fairness criterion could be implemented in a large-scale network. In this thesis, we show that simple rate control algorithms, using additive increase/multiplicative decrease rules [22] or explicit rates based on resource shadow prices, can provide stable convergence to *weighted proportional fairness*.

A vector x is such that the rates per unit charge are weighted proportionally fair for a weight vector w, if x is feasible, and if for any other feasible vector x^*

$$\sum_{r \in R} w_r \frac{x_{r^*} - x_r}{x_r} \leq 0. \tag{1.2}$$

The relationship between the conditions (1.1)–(1.2) is well illustrated when $w_r, r \in R$, are all integral. For each $r \in R$, replace the single user r by w_r identical sub-servers, construct the proportionally fair allocation over the resulting $\sum_r w_r$ users, and provide to user r the aggregate rate allocated to its w_r sub-users; then the resulting rates per unit charge are proportionally fair. Weighted proportional sharing appears as a means to achieve an allocation with optimal utility when users dynamically express the utility they attach to an allocation through the value they attribute to the weight parameters. (This construction also illustrates the need to adopt the notion of fairness to a non-cooperative context, where it is possible for a single user to represent itself as several distinct users.) Roberts and Massoulié [88] argue that x_r is approximately proportional to w_r only when the number of routes sharing a link is large and the individual allocations are small.

Mechanisms by which supply and demand reach equilibrium have, of course, long been a central concern of economists, and there exists a substantial body of theory on the stability of what are termed *tatonnement processes* ([3, 50, 94]). From this point of view, the rate control algorithms described in this thesis are particular embodiments of a *Walrasian auctioneer*, searching for market clearing prices. The Walrasian auctioneer of tatonnement theory is usually considered a rather implausible construct; we show that the structured nature of a communication network provides a very natural context within which we can investigate the consequences for a tatonnement process of stochastic perturbations and time-lags.

There is a substantial literature on rate control algorithms, recently reviewed by Hernandez-Valencia et al. [52]. Key early papers of Jacobson [57] and Chiu and Jain [22] identified the advantages of adaptive schemes that either increase flows linearly or decrease flows multiplicatively, depending on the absence or presence of congestion. Important recent papers of Bolot and Shankar [15], Fendick et al. [32] and Bonomi et al. [16] have analysed the stability of networks with a single bottleneck resource, where congestion is signalled by the build-up of a queue at the bottleneck's buffer, and where propagation delays are significant. Bolot and Shankar [15] modelled one and two connections involving a single bottleneck, obtaining closed-form solutions for the one connection case in both transient and steady-state regime. They also discussed the need for a trade-off between good steady-state behaviour and rapid convergence to equilibrium. A class of feedback schemes was analysed by Fendick et al. [32] using a system of retarded differential equations [92]. Bonomi et al. [16] discussed the importance of gain and damping parameters for stabilising behaviour and achieve fairness.

The framework we adopt in this thesis is simpler than that analysed by these authors in that we directly model only rates and not queue lengths, but more complex in that we model a network with an arbitrary number of bottleneck resources. Theoretical work ([27], [65]) on queues serving the superposition of a large number of streams indicates circumstances when the busy period preceding a buffer overflow may be relatively short. Several authors have argued the advantages of preventing queue build-up through the bounding of rates (see, for example, Charny et al. [21]). We also model two important aspects of decentralised systems, namely stochastic perturbations and time-lags. Both these factors provide important insights into the behaviour of the network, and allow us to quantify the various relationships and trade-offs between speed of convergence, the magnitude of fluctuations about the equilibrium point, and the stability of the network.

The organisation of the dissertation is as follows. In the next chapter, we apply the idea of a smart market [76] to develop an algorithm for a network protocol, combining a market mechanism with the additive increase/multiplicative decrease properties of TCP. We start with the system optimisation problem [66] and its decomposition into the user and the network subproblems. The optimisation approach to flow control has generated a series of works to steer the network towards optimality [5, 48, 72, 73, 74]. We then analyse two classes of rate control algorithms in terms of stability and rates of convergence. Stability is established by showing that the network's implicit objective function provides a Lyapunov function [11, 53, 62] for the dynamical system defined by the rate control algorithm. Similar approaches have been used by Mo and Walrand [83], Low and Lapsley [74], and Golestani and Bhattacharya [48]. Mo and Walrand studied the convergence of a fair end-to-end window-based congestion control protocol for a multicast fluid model of the network which generalises proportional fairness and includes arbitrarily close approximations of max-min fairness. In the scheme proposed by Low and Lapsley, users select transmission rates that maximise their own benefits, utility minus bandwidth cost, and network resources adjust bandwidth prices to coordinate the users' decisions. For Golestani and Bhattacharya [48], the proposal uses a minimum cost maximal flow algorithm. Massoulié and Roberts [77] discusses the design of distributed algorithms for sharing network bandwidth by comparing the criteria of max-min fairness, proportional fairness and minimum potential delay. In the second chapter, we also model two important aspects of decentralised systems, namely stochastic perturbations and time-lags.

A basic question for a network simulation is what topology to use for the network being simulated. Paxson and Floyd [86] argue that there is no such thing as a 'typical' Internet topology because it is constantly changing, and the topology is engineered by a number of different entities. We overcome this difficulty in Chapter 3 by analysing the rate control algorithms developed in Chapter 2, by considering a large random network. Since in our study we represent the topology of the network by an incidence matrix, this matrix will be randomly generated. Therefore, there is some regularity between the study of random matrices and our analysis. There is a tremendous literature on random matrices surveyed in Mehta [80] and Carmeli [18]. In the same chapter, we present an eigenvalue analysis of the rate control algorithms and deduce some interesting relationships between the model parameters and trade-offs among convergence rates, magnitude of oscillations about the stable point, and the stability of the network. It is worthwhile mentioning that the eigenvalues of the linearised systems satisfy the conditions for the law of large numbers for random matrices to hold. We end this chapter with an analysis of a network with a definite structure. The results show that there exists some regularity between different networks irrespective of their structures.

In Chapter 4, we investigate the stability of a simple network having a single bottleneck resource. This study will help us to understand the microscopic behaviour of a variant of the rate control algorithm introduced in Chapter 2 and will explain some aspects overlooked in the previous chapter. We consider a TCP-like algorithm and analyse the relationship between the gain parameters and the round-trip times for each user. The interesting feature of the algorithm in this chapter is the packet marking mechanism as congestion indicators. The premise of this approach is that if the resource implications of their actions can be made known to end-nodes, then the end-nodes themselves are best placed to determine what should be their demands upon the resources of the network. A body of work is emerging in this direction. Floyd [33] introduced the Explicit Congestion Notification for the TCP. Our work in this chapter is more closely related to that of Gibbens and Kelly [44, 45]. For a network with cooperative end-nodes, the marks provide the information the end-nodes need to make efficient use of the network. The relationship between time-lags and stability has been the focus of study of many researchers. Recently, Chong and his co-authors [23] proposed a First-order Rate-based Flow Control (FRFC) where he derived a sufficient condition for the stability of a single bottleneck resource with multiple connections of diverse round-trip delays. This issue was further studied by Tan [93] and Johari and Tan [60]. In this chapter, we explore this aspect for our model and show how the end-users may select the gain parameters, assuming the round-trip times of their respective routes to be known, in order to keep the system stable while using a fair share of the bandwidth.

Chapter 5 illustrates the theoretical results with various simulation examples. Extensive simulations of the algorithms described in this dissertation were performed in order to illustrate, qualitatively and quantitatively, some of the issues in the presentation. We finally conclude with some remarks and open issues.

Chapter 2

Network Optimisation

Optimisation based flow control is motivated by a welfare maximisation problem and derives its control mechanism as a distributed algorithm to solve it. An important feature of this approach is that sources with different valuation of bandwidth should react differently to congestion. Faced with the same congestion, a source with a higher valuation should send at a higher rate that one with a lower valuation [72]. This has the advantage of allowing the users to decide their payments and receive what the network allocates. Though it may not be possible, nor critical, that optimality is exactly attained in a real network, the optimisation framework offers a means to explicitly drive the network *as a whole* towards a desirable operating point.

The purpose of this chapter is to give a coherent description of an optimisation approach to flow control. We decompose the welfare maximisation problem into user and network optimisation subproblems, where a user specifies a charge and the network searches for a rate vector such that the rate per unit charge are proportionally fair (defined in Chapter 1). The optimisation problem results in control mechanisms to be implemented at the sources and the network resources. The algorithms adapt to fluctuations in the capacity of and the demand for these resources. They consist of two components: the *link algorithm* executed by network devices such as routers or switches, and the *source algorithm* executed by host computers or edge devices. The link algorithm detects congestion and feeds back information to users, and in response, the source algorithm adjusts the rate into the network. The basic design issue is what to feed back (link algorithm) and how to react (source algorithm) and the goal is to achieve stability, optimality and fairness. A sample of papers which describe the implementation of the source algorithms on TCP Tahoe, Reno and Vegas for the Internet are [17, 30, 82].

Several important schemes on the Internet do not have a link algorithm and require the source algorithm to infer network congestion from local observation such as round-trip times and packet loss. However, a new approach called Random Early Marking (REM) [4, 6, 71] does implement the link algorithm. There is a substantial literature on rate control algorithms, recently reviewed by Hernandez-Valencia et al. [52]. The algorithms we consider in this dissertation provide natural generalisations to large-scale networks of simple additive increase/multiplicative decrease schemes. Key early papers of Jacobson [57] and Chiu and Jain [22] identified the advantages of adaptive schemes that either increase flows linearly or decrease flows multiplicatively, depending on the absence or presence of congestion.

In the next section, we describe the basic model of a network. We present the welfare optimisation problem $SYSTEM(U, A, C)$ and its decomposition into the user and the network optimisation subproblems $NETWORK(A, C; w)$ and $DUAL(A, C; w)$ respectively. In Sections 2.2 and 2.3, we study two classes of decentralised algorithms, the primal and the dual algorithms, designed to implement solutions to relaxations of the corresponding subproblems. We prove their global stability under mild conditions on the utility functions and determine their rates of convergence. The stochastic and time-lag models representing a more realistic description of the system are also discussed. We end this chapter with some variants of the distributed algorithms.

2.1 The Optimisation Framework

We consider a network with a set J of *resources*, and let C_j be the finite *capacity* of resource j, for $j \in J$. Let a *source* r be a non-empty subset of J, and write R for the set of possible *routes*. Set $A_{jr} = 1$ if $j \in r$, so that resource j lies on route r, and set $A_{jr} = 0$ otherwise. This defines a $0-1$ matrix $A = (A_{jr}, j \in J, r \in R)$.

We assume that several routes through the network may substitute for one another and that each *source-sink* is served by a single route r, $r \in R$. We associate a source-sink r with a user, and suppose that if a rate x_r is allocated to the source-sink r then this has utility $U_r(x_r)$ to the user. In this thesis, we shall assume that the utility $U_r(x_r)$ is an increasing, strictly concave and continuously differentiable function of x_r over the range $x_r \geq 0$. A traffic that leads to such a utility function is termed *elastic* [89]. Assume further that $U_r(x_r)$ is continuously differentiable with $U_r'(x_r) \rightarrow \infty$ as $x_r \downarrow 0$ and $U_r'(x_r) \rightarrow 0$ as $x_r \uparrow \infty$. Moreover, we also assume that the utilities are additive, so that the aggregate utility of rates $x = (x_r, r \in R)$ is $\sum_{r\in R} U_r(x_r)$. Let $U = (U_r(.), r \in R)$. The objective is to choose source rates x so as to maximise the aggregate utility. Under this model, the system optimal rates solve the following problem.

$SYSTEM(U, A, C)$**:**

$$\text{maximise} \quad \sum_{r\in R} U_r(x_r) \tag{2.1}$$

$$\text{subject to} \quad Ax \leq C \tag{2.2}$$

$$\text{over} \quad x \geq 0. \tag{2.3}$$

The objective function in the above problem is differentiable and strictly concave with a compact feasible region; hence a unique maximising value of the vector x exists and can be found by Lagrangian methods [98]. However, it is impractical to solve this optimisation problem since it involves utilities U that are unlikely to be known by the network. This leads us to consider a decomposition of the welfare maximisation problem $SYSTEM(U, A, C)$ into two simpler problems; a user subproblem and a network subproblem.[1]

Suppose that the user r, $r \in R$, may choose an amount to pay per unit time, w_r, and receives in return a flow x_r proportional to w_r, say $x_r = w_r/\lambda_r$, where λ_r could be

[1] A new approach to solve the welfare maximisation problem is to minimise its dual [74].

regarded as a charge per unit flow for user r. Then the utility maximisation problem for user r is as follows.

$USER_r(U_r; \lambda_r)$**:**

$$\text{maximise} \quad U_r\left(\frac{w_r}{\lambda_r}\right) - w_r \tag{2.4}$$
$$\text{over} \quad w_r \geq 0. \tag{2.5}$$

Next, suppose that the network knows the vector $w = (w_r, r \in R)$ and attempts to maximise the function $\sum_r w_r \log x_r$. The network's optimisation problem is then as follows.

$NETWORK(A, C; w)$**:**

$$\text{maximise} \quad \sum_{r \in R} w_r \log x_r \tag{2.6}$$
$$\text{subject to} \quad Ax \leq C \tag{2.7}$$
$$\text{over} \quad x \geq 0. \tag{2.8}$$

In [66], Kelly has proved that there always exist prices $\lambda = (\lambda_r, r \in R)$, charges $w = (w_r, r \in R)$ and source rates $x = (x_r, r \in R)$, satisfying $w_r = \lambda_r x_r$ for $r \in R$, such that w_r solves the subproblem $USER_r(U_r; \lambda_r)$ for r belonging to the set R and x solves the subproblem $NETWORK(A, C; w)$; moreover, the vector x is the unique solution to the welfare maximisation problem $SYSTEM(U, A, C)$. From the convexity of the feasible region for x and the strict concavity of the objective function, and for the same $w_r = 1$, $r \in R$, a vector of rates x solves the subproblem $NETWORK(A, C; w)$ if and only if the rates per unit charge are proportionally fair (also the Nash bargaining solution [78]).

We note that if, for a fixed set of users and arbitrary parameters $w = (w_r, r \in R)$, the network solves $NETWORK(A, C; w)$, then, the resulting rates $x = (x_r, r \in R)$ solve a variant of the problem $SYSTEM(U, A, C)$, with a weighted objective function

$\sum_r \alpha_r U_r(x_r)$ where $\alpha_r = w_r/(x_r U_r'(x_r))$ for $r \in R$. Thus, a choice of parameters $w = (w_r, r \in R)$ by the network (rather than by the users) corresponds to an implicit weighting by the network of the relative utilities of different users, with weights related to the users' various marginal utilities.

It is worth mentioning that the logarithmic utility functions are ultimately associated with the concept of proportional fairness. Thus, if each user is able to choose a charge per unit time that he is able to pay, and if the rates per unit charge are proportionally fair, then a system optimum is achieved when users' choices of charges and the network's choice of allocated rates are in equilibrium. Under the decomposition of the problem $SYSTEM(U, A, C)$ into the problems $NETWORK(A, C; w)$ and $USER_r(U_r; \lambda_r)$, $r \in R$, the utility function $U_r(x_r)$ is not required by the network, and only appears in the optimisation problem faced by user r. It is therefore sufficient to study the problem $NETWORK(A, C; w)$ to understand the behaviour of the communication system. In the sequel, we shall call the problem $NETWORK(A, C; w)$ as the *primal problem*.

The Lagrangian [98] for the problem $NETWORK(A, C; w)$ is

$$L(x, z; \mu) = \sum_{r \in R} w_r \log x_r + \mu^T (C - Ax - z)$$

where $z \geq 0$ is a vector of slack variables and μ is a vector of Lagrange multipliers (or *shadow prices*).

Observe that

$$\frac{\partial L}{\partial x_r} = \frac{w_r}{x_r} - \sum_{j \in J} A_{jr} \mu_j, \; r \in R,$$

and thus the unique optimum to the primal problem is given by

$$x_r = \frac{w_r}{\sum_{j \in J} A_{jr} \mu_j}, \; r \in R, \tag{2.9}$$

where $(x_r, r \in R)$, $(\mu_j, j \in J)$ solve

$$\mu \geq 0, \quad Ax \leq C, \quad \mu^T(C - Ax) = 0 \tag{2.10}$$

and relation (2.9).

To understand the behaviour of the communication system we can equivalently study the *dual problem* associated to the problem $NETWORK(A, C; w)$. This problem may be derived as follows. Maximising the Lagrangian over x and z gives

$$\begin{aligned} L(x, z; \mu) &= \sum_{r \in R} w_r \log\left(\frac{w_r}{\sum_{j \in J} A_{jr}\mu_j}\right) + \mu^T C - \mu^T Ax \\ &= \sum_{r \in R} w_r \log w_r - \sum_{r \in R} w_r \log\left(\sum_{j \in J} A_{jr}\mu_j\right) + \sum_{j \in J} A_{jr}\mu_j C_j - \sum_{r \in R} w_r. \end{aligned}$$

After removing the terms not dependent on the shadow prices μ, and converting it to a maximisation problem, we obtain the dual as follows.
$DUAL(A, C; w)$:

$$\text{maximise} \quad \sum_{r \in R} w_r \log\left(\sum_{j \in r} \mu_j\right) - \sum_{j \in J} \mu_j C_j \tag{2.11}$$

$$\text{over} \quad \mu \geq 0. \tag{2.12}$$

Given the users' *willingness-to-pay* w, the problem $DUAL(A, C; w)$ determines the vector μ which will in turn give the equilibrium rate vector x using (2.9). While the problems $NETWORK(A, C; w)$ and $DUAL(A, C; w)$ are mathematically tractable, it would be difficult to implement a solution in any centralised manner. A centralised processor, even if it were itself completely reliable and could cope with the complexity of the computational task involved, would have its lines of communication through the network vulnerable to delays and failures: see [74] for detailed explanations on these situations. Instead, interest focuses on algorithms which are decentralised and of a simple form: the challenge is to understand how such algorithms can be designed so that the network as a whole reacts intelligently to perturbations.

In the sections that follow, we describe two simple classes of decentralised algorithm, designed to implement solutions to relaxations of the subproblems $NETWORK(A, C; w)$ and $DUAL(A, C; w)$.

2.2 The Primal Rate Control Algorithm

Consider the system of differential equations

$$\frac{d}{dt} x_r(t) = \kappa_r \Big(w_r - x_r(t) \sum_{j \in J} A_{jr} \mu_j(t) \Big), \ r \in R, \tag{2.13}$$

where

$$\mu_j(t) = p_j \Big(\sum_{s \in R} A_{js} x_s(t) \Big). \tag{2.14}$$

From now onwards we assume that, unless otherwise specified, r ranges over the set R and j ranges over the set J. Also $\kappa_r > 0$. In the system (2.13)–(2.14), $\mu_j(t)$ is a congestion signal (price) generated by link j as a function of aggregate source rate at link j. The price faced by source r is the sum of the link prices in the path, $\sum_{j \in J} A_{jr} \mu_j(t)$. We shall call the problem (2.13)–(2.14) the *primal algorithm* or the source algorithm corresponding to the primal problem $NETWORK(A, C; w)$ described in Section 2.1. It is executed by the host computers or edge devices and adjusts the rate at which user traffics are injected into the network in response to the feedback received.

There are several ways to motivate the relations (2.13)–(2.14). For example, suppose that $p_j(y)$ is a price charged by resource j, per unit flow through resource j, when the total flow through resource j is y. Then, by adjusting the flow $x_r(t)$ on route r, in accordance with equations (2.13)–(2.14), the network attempts to equalise the aggregate cost of this flow, $x_r(t) \sum_{j \in J} A_{jr} \mu_j(t)$, with a target value w_r, for every user r in R. (For an enlightening description of the technological implementation of such algorithms in an ATM network, see Courcoubetis et al. [26]. The reader is also referred to [28] for implementing an Internet with differentiated services using proportional fairness.)

For an alternative motivation, suppose that resource j generates a continuous stream of feedback signals at a rate $p_j(y)$ when the total flow through resource j is $y = \sum_{s \in R} A_{js} x_s$. Suppose further that when resource j generates a feedback signal, a copy is sent to each user r whose route passes through resource j, where it is interpreted as a congestion indicator requiring some reduction in the flow x_r. Then equation (2.13) corresponds to a response by user r that comprises two components: a steady (additive) increase at rate proportional to w_r, and a multiplicative decrease at rate proportional to the stream of feedback signals received. See Chiu and Jain [22] and Jacobson [57] for schemes of these types; Hernandez-Valencia et al. [52] review several algorithms based on congestion indication feedback. This rate adjustment (2.13) is adopted in the same way as TCP Tahoe and Reno do for the IP networks [17, 72, 82]. Gibbens and Kelly [45] have proposed a marking scheme to implement the primal algorithm (2.13)–(2.14). In their scheme, the marks convey to a source the charge $x_r(t) \sum_{j \in J} \mu_j(t)$ and the source adjusts its rate to equalise the charge with its worth w_r.

We shall next prove the global stability of the system (2.13)–(2.14) *via* a properly defined Lyapunov function[2].We then determine the rate of convergence, and, by considering perturbations about the stable point, model stochastic effects of the system. We complete the section with a study of the time-lag version of system (2.13)–(2.14) when the gain parameters are fixed for all the users.

[2]Let $\bar{x}$ be an equilibrium for the dynamical system $x' = f(x)$. Let $V : U \longrightarrow \Re$ be a continuous function defined on a neighbourhood U of x, differentiable on $U - \bar{x}$, such that *(a)* $V(\bar{x}) = 0$ and $V(x) > 0$ if $x \neq \bar{x}$, *(b)* $\dot{V} \leq 0$ in $U - \bar{x}$. Then V is called a Lyapunov function for $\bar{x}$. The reader is referred to any of [39, 81, 97, 99, 102] for more details on Lyapunov functions.

2.2.1 Global stability

Consider the expression

$$\mathcal{U}(x) = \sum_{r\in R} w_r \log x_r - \sum_{j\in J} \int_0^{\sum_{s\in R} A_{js}x_s} p_j(y)dy \tag{2.15}$$

where $w_r > 0$, $r \in R$. We suppose that, for $j \in J$, the function $p_j(y)$, $y \geq 0$, is a non-negative, continuous, increasing function of y, not identically zero.

Theorem 2.1 *The strictly concave function $\mathcal{U}(x)$ is a Lyapunovfunction for the system of differential equations (2.13)–(2.14). The unique value x maximising $\mathcal{U}(x)$ is a stable point of the system, to which all trajectories converge.*

Proof. The assumptions on $w_r > 0$, $r \in R$ and $p_j(.)$, $j \in J$, ensure that $\mathcal{U}(x)$ is strictly concave on $x \geq 0$ with an interior maximum; the maximising value of x is thus unique. It is sufficient to show that $\dot{\mathcal{U}} = 0$ at the maximum only and $\dot{\mathcal{U}} > 0$ for all other x.
Observe that

$$\frac{\partial}{\partial x_r}\mathcal{U}(x) = \frac{w_r}{x_r} - \sum_{j\in J} A_{jr}p_j\Big(\sum_{s\in R} A_{js}x_s\Big); \tag{2.16}$$

setting these derivatives to zero identifies the maximum, which is of the form given by equation (2.9). Further

$$\begin{aligned}\frac{d}{dt}\mathcal{U}\big(x(t)\big) &= \sum_{r\in R} \frac{\partial \mathcal{U}}{\partial x_r}\cdot\frac{d}{dt}x_r(t)\\ &= \sum_{r\in R}\frac{\kappa_r}{x_r(t)}\bigg(w_r - x_r(t)\sum_{j\in J} A_{jr}p_j\Big(\sum_{s\in R} A_{js}x_s(t)\Big)\bigg)^2,\end{aligned}$$

establishing that $\mathcal{U}(x(t))$ is strictly increasing in t, unless $x(t) = x$, the unique x maximising $\mathcal{U}(x)$. The function $\mathcal{U}(x)$ is thus a Lyapunov function for the system (2.13)–(2.14), and the theorem follows ([47], Chapter 2).

■

We observe that the theorem does not rely upon the property that p_j, $j \in J$, are increasing functions, but only upon the weaker property that $\mathcal{U}(x)$ is strictly concave with an interior maximum.

The functions p_j, $j \in J$, may be chosen so that the maximisation of the Lyapunov function $\mathcal{U}(x)$ arbitrarily closely approximates the optimisation problem $NETWORK(A, C; w)$, and, in this sense, is a relaxation of the network problem. Define the continuous functions $p_j(y) = (y - C_j + \varepsilon)^+/\varepsilon^2$ for $j \in J$, and some real small positive quantity ε. Then, as $\varepsilon \to 0$, the maximisation of the Lyapunov function $\mathcal{U}(x)$ approximates arbitrarily closely the primal problem $NETWORK(A, C; w)$ defined by equations (2.7)–(2.8); in particular the vector x maximising $\mathcal{U}(x)$ approaches the solution x to relations (2.9) and (2.10). The derivative $p_j'(y)$ may however become arbitrarily large as the approximation improves.

2.2.2 Rate of convergence

We have seen, in Theorem 2.2.1, that the system (2.13)–(2.14) converges to a unique stable point. In this section, we investigate the rate of convergence, by linearisation about the stable point.

Let x identify the unique vector maximising $\mathcal{U}(x)$, let $\mu_j = p_j(\sum_{s:j\in s} x_s)$, and suppose p_j is differentiable at this point, with derivative p_j'. Let $x_r(t) = x_r + \kappa_r^{1/2} x_r^{1/2} y_r(t)$. Linearising the system (2.13)–(2.14) about x, we obtain, for all $r \in R$,

$$\begin{aligned}
\frac{d}{dt} y_r(t) &= -\left(\kappa_r y_r(t) \sum_{j\in J} A_{jr}\mu_j + \kappa_r^{1/2} x_r^{1/2} \sum_{j\in J} A_{jr} p_j' \sum_{s\in R} A_{js} \kappa_s^{1/2} x_s^{1/2} y_s(t)\right) \\
&= -\left(\kappa_r \frac{w_r}{x_r} y_r(t) + \kappa_r^{1/2} x_r^{1/2} \sum_j \sum_s p_j' A_{jr} A_{js} \kappa_s^{1/2} x_s^{1/2} y_s(t)\right).
\end{aligned}$$

Let $\mathcal{K} = \mathrm{diag}(\kappa_r,\, r \in R)$. We may then write the system of equations in matrix form as

$$\frac{d}{dt}y(t) = -\left(\mathcal{K}WX^{-1} + \mathcal{K}^{1/2}X^{1/2}A^TP'AX^{1/2}\mathcal{K}^{1/2}\right)y(t) \tag{2.17}$$

where $y(t) = (y_r(t), r \in R)$, $X = \text{diag}(x_r, r \in R)$, $W = \text{diag}(w_r, r \in R)$ and $P' = \text{diag}(p'_j, j \in J)$ respectively.
Let

$$\Gamma^T\Phi\Gamma = \mathcal{K}WX^{-1} + \mathcal{K}^{1/2}X^{1/2}A^TP'AX^{1/2}\mathcal{K}^{1/2} \tag{2.18}$$

where Γ is an orthogonal matrix, $\Gamma^T\Gamma = I$, and $\Phi = \text{diag}(\phi_r, r \in R)$ is the matrix of eigenvalues, necessarily positive, of the real, symmetric, positive definite matrix (2.18). Then

$$\frac{d}{dt}y(t) = -\Gamma^T\Phi\Gamma y(t). \tag{2.19}$$

In this case, the gain parameter is implicitly involved in the system. The rate of convergence of this system to equilibrium is determined by the smallest eigenvalue, ϕ_r, of the matrix (2.18). Although the speed of convergence increases with the magnitudes of the derivatives P', it also depends on the gain parameters of the individual routes, denoted by the matrix $\mathcal{K}$. We shall see in Sections 2.2.3 and 2.2.4 that this conclusion requires qualification in the presence of either stochastic effects or of time-lags.

Our early assumption that $p_j(y)$, $j \in J$, are increasing functions is convenient and often natural: it implies that $\mathcal{U}(x)$ is strictly concave with an interior maximum. If the functions $p_j(y)$, $j \in J$, are not increasing, then situations may exist where $\mathcal{U}(x)$ may not have an interior maximum or it may have multiple stationary points. Provided the $p_j(y)$, are differentiable at a stationary point, the local behaviour near the stationary point is described by the equation (2.19).

The Lyapunov function (2.15) thus provides an enlightening analysis of the global stability of the system (2.13)–(2.14), and of the relationship between this system and the problem $NETWORK(A, C; w)$.

Remark 2.1 *Suppose that the gain parameters are assumed to be fixed, say κ, for all the users. Then the system (2.13)–(2.14) reduces to*

$$\frac{d}{dt} x_r(t) = \kappa_r \Big(w_r - x_r(t) \sum_{j \in J} A_{jr} \mu_j(t) \Big), \; r \in R, \tag{2.20}$$

where

$$\mu_j(t) = p_j \Big(\sum_{s \in R} A_{js} x_s(t) \Big). \tag{2.21}$$

The system (2.20)–(2.21) can be motivated in the same way as for the system (2.13)–(2.14), and the Lyapunov function $\mathcal{U}(x)$, defined by equation (2.15), still holds.

By linearising the system (2.20)–(2.21) about the stable point x, we have

$$\frac{d}{dt} y(t) = -\kappa \Gamma^T \Phi \Gamma y(t), \tag{2.22}$$

where

$$\Gamma^T \Phi \Gamma = W X^{-1} + X^{1/2} A^T P' A X^{1/2}. \tag{2.23}$$

The above representation of the dynamical system shows clearly the dependence of the convergence rate on the gain parameter, and will be used for a major part of our present investigation. Thus, from equation (2.22), we observe that the speed of convergence increases both with the gain parameter κ and with the magnitude of the derivatives P'.

The system (2.13)–(2.14) (or (2.20)–(2.21)) has omitted to model two important aspects of decentralised systems, namely stochastic perturbations, and time-lags. In the next section, we analyse the stochastic version of system (2.20)–(2.21) by considering small perturbations to the stable point x.

2.2.3 Stochastic analysis

Stochastic perturbations within the network may well arise from a resource's method of sensing its load. Equation (2.21) represents the response $\mu_j(t)$ of resource j, $j \in J$, as a continuous function of a load, $y = \sum_{s\in R} A_{js}x_s$, which is assumed known. In practice, a resource may assess its load by an error-prone measurement mechanism, and then choose a feedback signal from a small set of possible signals. Hernandez-Valencia et al. [52] and Bonomi et al. [16] provide more detailed descriptions of binary feedback and congestion indication rate control algorithms.

Consider a stochastic perturbation of the linearised equation (2.22). Let

$$dy(t) = -\kappa\Big(\Gamma^T\Phi\Gamma y(t)dt + FdB(t)\Big) \tag{2.24}$$

where F is an arbitrary $|R| \times |I|$ matrix and $B(t) = (B_i(t), i \in I)$ is a collection of independent standard Brownian motions, extended to $-\infty < t < \infty$. (Later in this section we describe how the modelling of different sources of randomness may lead to various explicit forms for the matrix F.)

The stationary solution to the system (2.24) is

$$y(t) = -\kappa \int_{-\infty}^{t} e^{-\kappa(t-\tau)\Gamma^T\Phi\Gamma} FdB(\tau), \tag{2.25}$$

or simply

$$y(t) = -\kappa \int_{-\infty}^{0} e^{-\kappa\tau\Gamma^T\Phi\Gamma} FdB(\tau), \tag{2.26}$$

by substituting τ for $(\tau - t)$. This can be checked by differentiating both sides of equation (2.25) with respect to t. This leads us back to (2.24) as required. The solution (2.25) is a linear transformation of the Gaussian process $(B(\tau), \tau < t)$; hence

$y(t)$ has a multivariate normal distribution, $y(t) \sim N(0, \Sigma)$, where

$$\begin{aligned}\Sigma &= \mathbb{E}\Big[y(t)y(t)^T\Big] \\ &= \kappa^2 \int_{-\infty}^{0} e^{\kappa\tau\Gamma^T\Phi\Gamma} F F^T e^{\kappa\tau\Gamma^T\Phi\Gamma} d\tau\end{aligned}$$

since the $B_j(t)$'s, $j \in J$ are independent. But

$$\begin{aligned}e^{\kappa\tau\Gamma^T\Phi\Gamma} &= I + \kappa\tau\Gamma^T\Phi\Gamma + \frac{\kappa^2\tau^2}{2!}\Gamma^T\Phi^2\Gamma + \dots \\ &= \Gamma^T e^{\kappa\tau\Phi}\Gamma,\end{aligned}$$

therefore substituting τ by $\kappa\tau$, we obtain

$$\Sigma = \kappa\Gamma^T \left[\int_{-\infty}^{0} e^{\tau\Phi}\Gamma F F^T \Gamma^T e^{\tau\Phi} d\tau\right] \Gamma.$$

But $e^{\tau\Phi} = \mathrm{diag}(e^{\tau\phi_i})$, where ϕ_i are the diagonal elements of the diagonal matrix Φ. Define the symmetric matrix $[\Gamma F; \Phi]$ by

$$\begin{aligned}[\Gamma F; \Phi]_{rs} &= \left[\int_{-\infty}^{0} e^{\tau\Phi}\Gamma F F^T \Gamma^T e^{\tau\Phi} d\tau\right]_{rs} \\ &= \int_{-\infty}^{0} e^{\tau\Phi}[\Gamma F F^T \Gamma^T]_{rs} e^{\tau\Phi} d\tau \\ &= \frac{[\Gamma F F^T \Gamma^T]_{rs}}{\phi_r + \phi_s}, \; r, s \in R\,. \qquad (2.27)\end{aligned}$$

Then

$$\Sigma = \kappa\Gamma^T[\Gamma F; \Phi]\Gamma. \qquad (2.28)$$

We interpret expression (2.28) as follows. The covariance matrix increases linearly with the gain parameter κ; as κ increases, the faster convergence to equilibrium described by relation (2.22) is at the cost of a greater spread at equilibrium. Varying the derivatives P' has a more subtle effect (through relation (2.23) and the construction of (2.28)) on the covariance matrix; broadly, as P' increases, not only is convergence to equilibrium faster, but also the spread at equilibrium decreases. However, we shall see in Section 2.2.4 that, in the presence of time-lags, increasing P' may compromise stability.

We next illustrate how various sources of randomness may lead to different covariance structures.

Congestion indication with joint feedback

Consider the following stochastic version of the system (2.20)–(2.21). Let $(N_j(\tau), \tau \geq 0)$, for $j \in J$, be a collection of independent unit rate Poisson processes, and let

$$dx_r(t) = \kappa \left(w_r dt - x_r(t) \sum_{j \in J} A_{jr} \varepsilon_j dN_j \left(\varepsilon_j^{-1} \int_0^t \mu_j(\tau) d\tau \right) \right) \tag{2.29}$$

where the functions $\mu_j(\cdot)$, for $j \in J$, are given by equation (2.21) and ε_j are small real positive quantities. Equation (2.29) would describe the following model: resource j generates feedback signals indicating congestion as a time-dependent Poisson process at rate $\varepsilon_j^{-1}\mu_j(t)$; when resource j generates a feedback signal, a copy is sent to each user r whose route passes through resource j; and user r reacts to such a feedback signal by reducing $x_r(t)$ by an amount $\kappa\varepsilon_j x_r(t)$.

Now as $\varepsilon_j \to 0$, the normalised Poisson process $((\varepsilon_j N_j(\tau/\varepsilon_j) - \tau)\varepsilon_j^{-1/2}, \tau \geq 0)$ converges in distribution to a standard Brownian motion [14]. This motivates the approximation, valid when ε_j are small, of the Poisson driving equation (2.29) by its Brownian version

$$dx_r(t) = \kappa \left(w_r dt - x_r(t) \sum_{j \in J} A_{jr} \left(\mu_j(t) dt + \varepsilon_j^{1/2} \mu_j(t)^{1/2} dB_j(t) \right) \right) \tag{2.30}$$

where $(B_j(t), t \geq 0)$, for $j \in J$, are a collection of independent standard Brownian motions, representing stochastic effects at resource j, and ε_j is a scaling parameter for these effects. If the scaling parameters $\varepsilon_j, j \in J$, are small then the stochastic differential equation (2.30) has, as solution, a multidimensional Ornstein–Uhlenbeck process [102], centred on the stable point x of the differential equations (2.20)–(2.21). The stationary distribution for $(x_r(t), r \in R)$ is a multivariate normal distribution, with covariance matrix that can be explicitly calculated in terms of the parameters of the network. The covariance matrix gives information about the complex coordination among potentially all sources due to coupling of sources through shared links.

The corresponding Brownian version of the linearised system (2.22) is simply equation (2.24) where $B(t) = (B_j(t), j \in J)$, and F is an $|R| \times |J|$ matrix with elements

$$F_{rj} = \varepsilon_j^{1/2} \mu_j^{1/2} A_{jr} x_r^{1/2}, \; r \in R, j \in J. \tag{2.31}$$

Then

$$FF^T = X^{1/2} A^T E P A X^{1/2}, \tag{2.32}$$

where $E = \mathrm{diag}(\varepsilon_j, j \in J)$ and $P = \mathrm{diag}(\mu_j, j \in J)$. Hence the stationary covariance matrix Σ may be calculated from expression (2.28). The scheme presented in this section is called *congestion indication with joint feedback.*

Congestion indication with individual feedback

Next, consider the Poisson driving equation

$$dx_r(t) = \kappa \left(w_r dt - \sum_{j \in r} \varepsilon_j dN_{jr} \left(\varepsilon_j^{-1} \int_0^t x_r(\tau) \mu_j(\tau) d\tau \right) \right) \tag{2.33}$$

where $(N_{jr}(\tau), \tau \geq 0)$, for $j \in J$, $r \in R$, are a collection of independent unit rate Poisson processes. This would describe the following model: feedback signals from resource j to user r arise at rate $\varepsilon_j^{-1} x_r(t) \mu_j(t)$; and user r reacts to such a feedback signal by reducing $x_r(t)$ by an amount $\kappa\varepsilon_j$. The Brownian approximation, valid when

ε_j, $j \in J$, are small is

$$dx_r(t) = \kappa\Bigg(w_r dt - x_r(t) \sum_{j \in J} A_{jr}\Big(\mu_j(t)dt + \varepsilon_j^{1/2} x_r(t)^{-1/2} \mu_j(t)^{1/2} dB_{jr}(t)\Big)\Bigg),$$

whose linearisation is equation (2.24) where the $|R| \times |J|.|R|$ matrix F is given by

$$F_{r,(j,s)} = \varepsilon_j^{1/2} \mu_j^{1/2} A_{jr} I[r = s]. \tag{2.34}$$

Thus FF^T is the matrix diag$(\sum_{j \in r} \mu_j \varepsilon_j, r \in R)$, and the stationary covariance matrix Σ may be calculated from expression (2.28). In Chapter 3 we provide a numerical illustration of this calculation, and contrast the results derived from expressions (2.31) and (2.34). The above type of feedback is called *congestion indication with individual feedback.*

Source fluctuations

Consider the Brownian driving equation

$$dx_r(t) = \kappa\Bigg(w_r dt - x_r(t) \sum_{j \in r} \mu_j(t)dt + \varepsilon_r^{1/2} x_r(t)^{1/2} dB_r(t)\Bigg).$$

This might correspond to fluctuations arising at sources, rather than within the network. In this case the $|R| \times |R|$ matrix F is the diagonal matrix diag$(\varepsilon_r^{1/2}, r \in R)$.

The results obtained by assuming the gain parameters to be route dependent do not indicate clearly the effects of these quantities on the covariance matrix since the gain parameters are implicitly involved.

2.2.4 Time-Lags

Next, we describe a model incorporating time-lags that generalises equations (2.20)–(2.21), and shall analyse its behaviour close to the stable point x. Our models of stochastic effects and of time-lags provide important insights into the behaviour of the network, and allow us to quantify the various relationships and trade-offs between speed of convergence, the magnitude of fluctuations about the equilibrium point, and the stability of the network.

Consider the lagged, discrete time system

$$x_r[t+1] = x_r[t] + \kappa\left(w_r - x_r[t]\sum_{j\in J} A_{jr}\mu_j[t-d(j,r)]\right) \tag{2.35}$$

where

$$\mu_j[t] = p_j\left(\sum_{s\in R} A_{js}x_s[t-d(j,s)]\right), \tag{2.36}$$

and $d(j,r)$, $j \in J$, $r \in R$, are non-negative integers. This might correspond to a model of congestion indication with joint feedback (Section 2.2.3), where there is a delay of $d(j,r)$ between resource j generating a feedback signal and user r receiving it, and another delay of $d(j,r)$ between user r changing its rate and the altered flow reaching resource j. Say that a vector x is an *equilibrium point* of the system (2.35)–(2.36) if $x_r[t] = x_r$, for $t = \dots, 0, 1, 2, \dots$, satisfies these equations. We use the same Lyapunov function $\mathcal{U}(x)$ defined by equation (2.15) in Section 2.2.1 and prove the following result.

Theorem 2.2 *The vector x maximising the strictly concave function $\mathcal{U}(x)$ is the unique equilibrium point of the system (2.35)–(2.36).*

Proof. The vector x is an equilibrium point if and only if it solves

$$w_r = x_r \sum_{j\in J} A_{jr}p_j\left(\sum_{s\in R} A_{js}x_s\right),\ j \in J,\ r \in R.$$

But this is precisely the stationary condition implied by the partial derivatives (2.16) of the function $\mathcal{U}(x)$, a strictly concave function with a unique maximum. This completes the proof.

■

For small enough values of κ the equilibrium point will be asymptotically stable, since if we replace κ by $\kappa\delta$ in equation (2.35) and let $x_r(t) = x_r[t/\delta]$, then as $\delta \to 0$, we may approximate arbitrarily closely a solution to equations (2.20)–(2.21). However, for small values of κ, convergence to the equilibrium point is slow, and so it is of interest to investigate the local stability of the equilibrium point for general values of κ.

We now determine the rate of convergence by considering perturbations about the stable point.
Let $\mu_j = p_j\left(\sum_{s\in R} A_{js}x_s\right)$, $j \in J$, and suppose p_j is differentiable at this point, with derivative p_j'. Let $x_r[t] = x_r + x_r^{1/2}y_r[t]$. Then, linearising the system (2.35)–(2.36) about x, we obtain

$$\begin{aligned} & y_r[t+1] \\ = \;& y_r[t] - \kappa\left(y_r[t]\sum_{j\in r}\mu_j + x_r^{1/2}\sum_{j\in r}p_j'\sum_{s\in R}A_{js}x_s^{1/2}y_s[t-d(j,r)-d(j,s)]\right) \\ = \;& y_r[t] - \kappa\left(\frac{w_r}{x_r}y_r[t] + \sum_{j\in J}\sum_{s\in R}p_j'A_{jr}A_{js}x_r^{1/2}x_s^{1/2}y_s[t-d(j,r)-d(j,s)]\right) \end{aligned} \tag{2.37}$$

for all $r \in R$. We proceed to express the system (2.37) in matrix form. Define the $|R| \times |R|$ matrices $(L[d], d = 0, 1, \ldots, D)$ by

$$(L[d])_{rs} = \sum_j p_j' A_{jr}A_{js}x_r^{1/2}x_s^{1/2}I[d(j,r)+d(j,s)=d]$$

where $D = \max_{j,r,s}\{d(j,r)+d(j,s)\}$ is the largest in the system and I is the indicator function. Thus

$$\sum_{d=0}^{D} L[d] = X^{1/2}A^T P' A X^{1/2},$$

the second term of the key matrix (2.23). Define the vector $y[t] = (y_r[t], r \in R)$.

Then, we can rewrite equation (2.37) in the matrix form

$$\begin{pmatrix} y[t+1] \\ y[t] \\ \vdots \\ y[t-D+1] \end{pmatrix} = L \begin{pmatrix} y[t] \\ y[t-1] \\ \vdots \\ y[t-D] \end{pmatrix} \tag{2.38}$$

where

$$L = \begin{pmatrix} I - \kappa(WX^{-1} + L[0]) & -\kappa L[1] & -\kappa L[2] & \dots & -\kappa L[D] \\ I & 0 & 0 & \dots & 0 \\ 0 & I & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}. \tag{2.39}$$

Observe that only the first row of the matrix L contains information about the system's behaviour. The equilibrium point x of the system (2.35)–(2.36) is stable if and only if the spectral radius of the matrix L is less than unity [12, 75]. Recall that in our model of stochastic effects (page 22), increasing the derivatives P' had broadly the same reductive effect on the covariance matrix Σ in (2.28) as *decreasing* the gain parameter κ; in contrast, the destabilising effect on the matrix L in (2.39) of increasing P' is broadly the same as *increasing* κ. This aspect is thoroughly investigated for a one-node system in Chapter 4.

For simplicity of notation, we have used the same gain parameter κ for each $r \in R$. If κ is replaced by κ_r in equation (2.35), then we again obtain relations (2.38)–(2.39), but now with κ interpreted as the matrix $\mathcal{K} = \operatorname{diag}(\kappa_r, r \in R)$. We investigate how the round-trip times within a network affect the choice of gain parameters in Chapter 4. As a continuation of this work, we may study the problem of choosing $\mathcal{K} = \operatorname{diag}(\kappa_r, r \in R)$ in order to minimise the spectral radius of matrix L.

There exist other natural discrete time versions of the equations (2.20)–(2.21), and these too may be analysed in a similar manner. For example, consider the method of repeated substitution, $x_r[t+1] = w_r / \sum_{j \in J} A_{jr}\mu_j[t - d(j,r)]$, $r \in R$, or its damped version

$$x_r[t+1] = (1-\kappa)x_r[t] + \kappa \frac{w_r}{\sum_{j \in J} A_{jr}\mu_j[t - d(j,r)]}$$

where $\mu_j[t]$ is again given by equation (2.36). Then, the linearised relations (2.38)–(2.39) are altered in that the top row of the matrix (2.39) becomes

$$\left(I - \kappa(I + XW^{-1}L[0]), -\kappa XW^{-1}L[1], \ldots, -\kappa XW^{-1}L[D]\right).$$

Next, we study the dual algorithm obtained from the problem $DUAL(A, C; w)$ in Section 2.1, page 15.

2.3 The Dual Algorithm

The equations (2.13)–(2.14) (or (2.20)–(2.21)) in Section 2.2 represent a system where rates vary gradually, and shadow prices are given as functions of the rates. Here, we investigate a system where shadow prices vary gradually, with rates given as functions of the shadow prices. We also assume the gain parameters to be fixed for all the end-users. The model when the gain parameters are resource dependent can be treated in the same way as for the primal rate control algorithm.

Let

$$\frac{d}{dt}\mu_j(t) = \kappa\left(\sum_{j \in J} A_{jr}x_r(t) - q_j\left(\mu_j(t)\right)\right) \tag{2.40}$$

where

$$x_r(t) = \frac{w_r}{\sum_{k \in J} A_{kr}\mu_k(t)} \tag{2.41}$$

and $\mu_j(t)$ is a congestion signal (price) generated by link j as a function of the aggregate source rate at link j. We shall call the system (2.40)–(2.41) as the *dual algorithm* or link algorithm. This algorithm is executed by network devices such as routers or switches. It detects congestion and feeds back information to users, and in response, the source algorithm adjusts the rate at which user traffics are injected into the network.

The relationship between the algorithm (2.40)–(2.41) and the problem $DUAL(A, C; w)$ (page 15) parallels that between the primal algorithm (2.13)–(2.14) and the problem $NETWORK(A, C; w)$ (page 13). As in Section 2.2 we may motivate the algorithm in several ways. For example, suppose that $q_j(\eta)$ is the flow through resource j which generates a price at resource j of η. An economist would describe the right hand side of equation (2.40) as the vector of excess demand at prices $(\mu_j(t), j \in J)$, and would recognise equations (2.40)–(2.41) as a tatonnement process by which prices adjust according to supply and demand (Varian [94], Chapter 21). Several important schemes in the Internet do not implement link algorithm. However, REM can be regarded as a marking implementation of the dual algorithm (2.40)–(2.41) with log utility functions: see [4, 6, 71, 72] for more details about REM. The approach of solving the dual of the main optimisation problem $SYSTEM(U, A, C)$ in Section 2.1 reduces to solving the system (2.40)–(2.41) under mild conditions.

In the following sections, we derive conditions for stability and rate of convergence and also consider the perturbation analysis of stochastic effects and time-lags.

2.3.1 Global stability

The global stability of the dual may be established using a similar approach to that in Section 2.2.1. Consider the expression

$$\mathcal{V}(\mu) = \sum_{r \in R} w_r \log\left(\sum_{j \in J} A_{jr}\mu_j\right) - \sum_{j \in J} \int_0^{\mu_j} q_j(\eta)d\eta \tag{2.42}$$

where $w_r > 0$, $r \in R$, and suppose that, for $j \in J$, $q_j(0) = 0$ and $q_j(\eta)$, $\eta \geq 0$, is a continuous, strictly increasing function of η.

Theorem 2.3 *The strictly concave function $\mathcal{V}(\mu)$ is a Lyapunov function for the system of differential equations (2.40)–(2.41). The unique value μ maximising $\mathcal{V}(\mu)$ is a stable point of the system, to which all trajectories converge.*

Proof. The assumptions on $w_r > 0$, $r \in R$ and on q_j, $j \in J$, ensure that $\mathcal{V}(\mu)$ given by (2.42) is strictly concave on $\mu \geq 0$ with an interior maximum; the maximising value of μ is thus unique, and is determined by setting the derivatives

$$\frac{\partial}{\partial \mu_j}\mathcal{V}(\mu) = \sum_{r \in R} \frac{A_{jr} w_r}{\sum_{k \in J} A_{kr}\mu_k} - q_j(\mu_j) \tag{2.43}$$

to zero. Also,

$$\begin{aligned} \frac{d}{dt}\mathcal{V}\big(\mu(t)\big) &= \sum_{j \in J} \frac{\partial \mathcal{V}}{\partial \mu_j} \cdot \frac{d}{dt}\mu_j(t) \\ &= \kappa \sum_{j \in J}\left(\sum_{r \in R} \frac{A_{jr} w_r}{\sum_{k \in J} A_{kr}\mu_k(t)} - q_j\big(\mu_j(t)\big)\right)^2, \end{aligned}$$

establishing that $\mathcal{V}(\mu(t))$ is strictly increasing with t, unless $\mu(t) = \mu$, the unique value μ maximising $\mathcal{V}(\mu)$. The function $\mathcal{V}(\mu)$ is thus a Lyapunov function for the system (2.40)–(2.41), and the theorem follows [47].

■

The maximisation of the Lyapunov function $\mathcal{V}(\mu)$ becomes the dual problem if, for $j \in J$, $\eta > 0$, $q_j(\eta) = C_j$. These functions violate the assumption that $q_j(\eta)$ is continuous at $\eta = 0$, but they may be arbitrarily closely approximated, for example by the functions

$$q_j(\eta) = C_j\eta/(\eta + \varepsilon)$$

for small positive ε. The derivative $q_j'(\eta)$ may however become arbitrarily large as the approximation improves.

Further, by appropriate choice of the functions q_j, $j \in J$, the maximisation of the function $\mathcal{V}(\mu)$ can arbitrarily closely approximate the problem $DUAL(A, C; w)$.

We investigate the stability of the system (2.40)–(2.41) by determining the rate of convergence about the stable point.

2.3.2 Rate of convergence

Suppose that μ identifies the unique vector maximising $\mathcal{V}(\mu)$ following Theorem 2.3.1. Let $x_r = w_r / \sum_{k \in r} \mu_k$, and suppose $q_j(\eta)$ is differentiable at the point $\eta = \mu_j$, with derivative q_j'. Further, assume $\mu_j(t) = \mu_j + \xi_j(t)$, $j \in J$. Then, linearising the system (2.40)–(2.41) about μ, we obtain, for all $j \in J$,

$$\begin{aligned}
\frac{d}{dt}\xi_j(t) &= \kappa\left(\sum_{r \in R} A_{jr} w_r \Big(\sum_{k \in J} A_{kr}(\mu_k + \xi_k(t))\Big)^{-1} - \Big(q_j + q_j'(\mu_j)\xi_j(t)\Big)\right) \\
&= \kappa\left(\sum_{r \in R} A_{jr} x_r \Big(1 - \frac{\sum_{k \in J} A_{kr}\xi_k(t)}{\sum_{k \in J} A_{kr}\mu_k}\Big) - \Big(q_j + q_j'(\mu_j)\xi_j(t)\Big)\right) \\
&= -\kappa\left(\sum_{r \in R} A_{jr} x_r^2 w_r^{-1} \sum_{k \in J} A_{kr}\xi_k(t) + q_j'(\mu_j)\xi_j(t)\right).
\end{aligned}$$

We may write the above equations in matrix form as

$$\frac{d}{dt}\xi(t) = -\kappa\left(AXW^{-1}XA^T + Q'\right)\xi(t) \tag{2.44}$$

where $\xi(t) = (\xi_j(t), j \in J)$, $X = \mathrm{diag}(x_r, r \in R)$, $W = \mathrm{diag}(w_r, r \in R)$ and $Q' = \mathrm{diag}(q_j', j \in J)$ respectively. Let

$$\Theta^T \Psi \Theta = AXW^{-1}XA^T + Q' \tag{2.45}$$

where Θ is an orthogonal matrix, $\Theta^T\Theta = I$, and $\Psi = \text{diag}(\psi_j, j \in J)$ is the matrix of eigenvalues, necessarily non-negative, of the real, symmetric, positive semidefinite matrix (2.45). Then

$$\frac{d}{dt}\xi(t) = -\kappa\Theta^T\Psi\Theta\xi(t), \tag{2.46}$$

and thus the rate of convergence to the stable point is determined by the smallest eigenvalue of the matrix (2.45). The speed of convergence increases both with the gain parameter κ and with the magnitude of the derivatives Q'. This conclusion requires qualification in the presence of either stochastic effects of time-lags.

We shall investigate the presence of stochastic effects on the linearised system (2.46).

2.3.3 Stochastic analysis

We consider a stochastic perturbation of the linearised equation (2.46). Let

$$d\xi(t) = -\kappa\left(\Theta^T\Psi\Theta\xi(t)dt - GdB(t)\right) \tag{2.47}$$

where $B(t) = (B_i(t), i \in I)$ is a collection of independent standard Brownian motions, and G is an arbitrary $|J| \times |I|$ matrix.

A similar analysis to that of Section 2.2.3 determines the stationary covariance matrix Σ of $\xi(t)$. The stationary solution of (2.47) is

$$\xi(t) = -\kappa\int_{-\infty}^{t} e^{-\kappa(t-\tau)\Theta^T\Psi\Theta}GdB(\tau), \tag{2.48}$$

as can be checked by differentiating both sides of (2.48) with respect to t. The solution (2.48) is a linear transformation of the Gaussian process $(B(\tau), \tau < t)$; hence $\xi(t)$ has a multivariate normal distribution, $\xi(t) \sim N(0, \Sigma)$, where Σ is obtained in the same way as that in the primal case, page 23.

$$\begin{aligned}\Sigma &= \mathbb{E}\Big[\xi(t)\xi(t)^T\Big] \\ &= \kappa\Theta^T\left[\int_{-\infty}^{0} e^{\tau\Psi}\Theta GG^T\Theta^T e^{\tau\Psi}d\tau\right]\Theta.\end{aligned}$$

By defining the symmetric matrix $[\Theta G;\Psi]$ by

$$[\Theta G;\Psi]_{jk} = \frac{[\Theta GG^T\Theta^T]_{jk}}{\psi_j+\psi_k},$$

we have

$$\Sigma = \kappa\Theta^T[\Theta G;\Psi]\Theta. \tag{2.49}$$

The covariance matrix increases linearly with the gain parameter κ; as κ increases, the faster convergence to equilibrium described by relation (2.46) is at the cost of a greater spread at equilibrium.

Motivated by this derivation we describe an example illustrating how a model of the form (2.47) might arise.

Shadow prices inferred from fluctuating flow rates

Consider the Poisson driving equation

$$d\mu_j(t) = \kappa\left(\sum_{r\in R} A_{jr}\varepsilon_r dN_r\left(\varepsilon_r^{-1}\int_0^t x_r(\tau)d\tau\right) - q_j\Big(\mu_j(t)\Big)dt\right)$$

where $(N_r(\tau),\tau\geq 0)$, for $r\in R$, are a collection of independent unit rate Poisson processes. This would describe a model where, on a very fine time-scale, the flow on route r takes the form of a time-dependent Poisson process of rate $x_r(t)/\varepsilon_r$, with each point of the process containing a workload of size ε_r. The Brownian approximation, valid for small positive ε_r, becomes

$$d\mu_j(t) = \kappa \Bigg(\sum_{r \in R} A_{jr} \Big(x_r(t)dt + x_r(t)^{1/2} \varepsilon_r^{1/2} dB_r(t) \Big) - q_j \Big(\mu_j(t) \Big) dt \Bigg) \tag{2.50}$$

where $B_r(t)$, $t \geq 0$, $r \in R$, is a standard Brownian motion, representing stochastic effects associated with the flow on route r. The linearisation of (2.50) is equation (2.47) where G is a $|J| \times |R|$ matrix with elements

$$G_{jr} = \varepsilon_r^{1/2} x_r^{1/2} A_{jr}; \tag{2.51}$$

thus $GG^T = AXEA^T$ where $E = \text{diag}(\varepsilon_r, r \in R)$.

If the scaling parameters ε_r, are small then the stationary distribution for $(\mu_j(t), j \in J)$ is centred on the stable point μ of the differential equations (2.40)–(2.41), with a covariance matrix that can be explicitly calculated in terms of the parameters of the network. The covariance matrix provides an insight about the coordination of the shadow prices for the various resources in the network.

2.3.4 Time-Lags

In this section, we study the discrete time-lag version of the dual algorithm (2.40)–(2.41), page 30, in a similar way to that of the primal algorithm in Section 2.2.4. Consider the system

$$\mu_j[t+1] = \mu_j[t] + \kappa \Big(\sum_{r \in R} A_{jr} x_r[t - d(j,r)] - q_j \left(\mu_j[t] \right) \Big) \tag{2.52}$$

where

$$x_r[t] = \frac{w_r}{\sum_{k \in J} A_{kr} \mu_k[t - d(k,r)]} \tag{2.53}$$

and $d(j,r)$, $j \in J$, $r \in R$, are non-negative integers. Suppose a vector μ is an equilibrium point of the system (2.52)–(2.53) if $\mu_j[t] = \mu_j$, for $t = \ldots, 0, 1, 2, \ldots$, satisfies these equations. We use the same Lyapunov function $\mathcal{V}$ defined in Section 2.3.1 and prove the following theorem.

Theorem 2.4 *The vector μ maximising the strictly concave function $\mathcal{V}(\mu)$ is the unique equilibrium point of the system (2.52)–(2.53).*

Proof. The vector μ is an equilibrium point if and only if it solves

$$\sum_{r \in R} \frac{A_{jr} w_r}{\sum_{k \in J} A_{kr} \mu_k} = q_j(\mu_j).$$

But this is precisely the stationary condition implied by the partial derivatives (2.43) of the function $\mathcal{V}(\mu)$, a strictly concave function with a unique maximum. The result follows.

■

We investigate the stability of the equilibrium point by linearising (2.52)–(2.53) about μ. Let $x_r = w_r / \sum_{k \in r} \mu_k$, and suppose q_j is differentiable at the point $y = \mu_j$, with derivative q_j'. Assume $\mu_j[t] = \mu_j + \xi_j[t]$. Linearising the system (2.52)–(2.53) about μ, we obtain

$$\xi_j[t+1] = \xi_j[t] - \kappa \Big(\sum_{r \in R} A_{jr} x_r^2 w_r^{-1} \sum_{k \in J} A_{kr} \xi_k \big[t - d(j,r) - d(k,r) \big] + q_j' \xi_j[t] \Big). \quad (2.54)$$

Define the $|J| \times |J|$ matrices $(M[d], d = 0, 1, \ldots, D)$ by

$$(M[d])_{jk} = \sum_{r \in R} x_r^2 w_r^{-1} A_{jr} A_{kr} I[d(j,r) + d(k,r) = d],$$

where in this case $D = \max_{j,k,r} \{ d(j,r) + d(k,r) \}$ and I is the indicator function. Thus

$$\sum_{d=0}^{D} M[d] = AXW^{-1}XA^T.$$

Define the vector $\xi[t] = (\xi_j[t], j \in J)$. We can then rewrite equation (2.54) in the matrix form as

$$\begin{pmatrix} \xi[t+1] \\ \xi[t] \\ \vdots \\ \xi[t-D+1] \end{pmatrix} = M \begin{pmatrix} \xi[t] \\ \xi[t-1] \\ \vdots \\ \xi[t-D] \end{pmatrix} \tag{2.55}$$

where

$$M = \begin{pmatrix} I - \kappa(M[0] + Q') & -\kappa M[1] & -\kappa M[2] & \dots & -\kappa M[D] \\ I & 0 & 0 & \dots & 0 \\ 0 & I & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & & \vdots \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}. \tag{2.56}$$

As in the time-lag version of the primal algorithm, only the first row of M carries the information about the system. The equilibrium point μ of the system (2.52)–(2.53) is stable if the spectral radius of the matrix M is less than unity [12, 75]. With stochastic effects, increasing the derivatives Q' has broadly the same reductive effect on the covariance matrix (2.49) as *decreasing* the gain parameter κ; in contrast, the destabilising effect on the matrix (2.56) of increasing Q' is broadly the same as *increasing* κ.

2.4 Variants

Several variants of the primal algorithm (2.13)–(2.14) (or (2.20)–(2.21)) and the dual algorithm (2.40)–(2.41) allow a similar analysis. For example, if the right hand side of equation (2.20) is multiplied by a positive function $f_r(x(t), \mu(t))$ then Theorem 2.2.1 remains valid. Similarly, if the right hand side of equation (2.40) is multiplied by a positive function $f_j(x(t), \mu(t))$ then Theorem 2.3.1 remains valid. As a simple example, we could divide the right hand side of equation (2.20) by w_r, or of equation (2.40) by $q_j(\mu_j(t))$.

Suppose equation (2.40) is replaced by

$$\frac{d}{dt}\mu_j(t) = \kappa\left(q_j\Big(\sum_{j\in J} A_{jr}x_r(t)\Big) - \mu_j(t)\right) \tag{2.57}$$

where

$$x_r(t) = \frac{w_r}{\sum_{k\in J} A_{kr}\mu_k(t)}. \tag{2.58}$$

and q_j is the inverse function of p_j (See also [68]). The above system (2.57)–(2.58) which is a variant of the dual algorithm (2.40)–(2.41) can be interpreted as follows: suppose that $q_j(y)$ is a price charged by resource j, per unit flow through resource j, when the total load on resource j is $\sum_{s\in R} A_{js}x_j$. Then by adjusting the (shadow) price $\mu_j(t)$ at resource j according to relations (2.57)–(2.58), the network attempts to equalise the aggregate cost per unit flow for all routes through the resource j, with target value w_r. Another interpretation can be in terms of rate of feedback generated by resource j.

The derivative of q_j, namely $q_j'(y)$, can be considered as the rate of change of the price of the resource when the total load on the resource is $y = \sum_{s\in R} A_{js}x_s$.

By following the same procedure described in Section 2.3, we obtain mathematically tractable forms for the rate of convergence of the system. We summarise the main results of this variant of the dual algorithm (2.40)–(2.41).

By linearising the system (2.57)–(2.58) about μ, where $\mu_j(t) = \mu_j + {q_j}'^{1/2}\xi_j(t)$, we obtain

$$\frac{d}{dt}\xi(t) = -\kappa(Q'^{1/2}AXW^{-1}XA^TQ'^{1/2} + I)\xi(t). \tag{2.59}$$

It is convenient to assume

$$\Theta^T \Psi \Theta = Q'^{1/2} A X W^{-1} X A^T Q'^{1/2} + I \tag{2.60}$$

where Θ is an orthogonal matrix, and $\Psi = \mathrm{diag}(\psi_j, j \in J)$ is the matrix of eigenvalues, necessarily positive, of the real, symmetric, positive definite matrix (2.60). Then

$$\frac{d}{dt}\xi(t) = -\kappa \Theta^T \Psi \Theta \xi(t). \tag{2.61}$$

Thus, the rate of convergence to the stable point is determined by the smallest eigenvalue of the matrix (2.60): the rate of convergence varies with the gain parameter κ and the matrix of derivatives Q'.

If we consider a stochastic perturbation of the linearised equation (2.61), we still observe that the covariance matrix increases linearly with the gain parameter κ in the way described in Section 2.3.3, page 34.

Further, suppose that the rates x_r may be approximated by independent Poisson processes N_r, giving rise to the following equation,

$$d\mu_j(t) = \kappa \left(q_j \left(\sum_{r \in R} A_{jr} \varepsilon_r dN_r \left(\varepsilon_r^{-1} \int_0^t x_r(\tau) d\tau \right) \right) - \mu_j(t) dt \right). \tag{2.62}$$

The above algorithm describes the model, where, on a very fine time-scale, the flow on route r takes the form of a time-dependent Poisson process of rate $x_r(t)/\varepsilon_r$, with each point of the process containing a workload of size ε_r. This in turn causes a fluctuation on the total load of resource j and consequently on the price charged by resource j, namely q_j.

Assuming that ε tends to zero, we obtain the Brownian approximation of the above equation as

$$d\mu_j(t) \quad = \quad \kappa \left(q_j \left[\sum_{r \in R} A_{jr} \left(x_r(t) dt + \varepsilon_r^{1/2} x_r(t)^{1/2} dB_r(t) \right) \right] - \mu_j(t) dt \right).$$

Then, linearising this system, we get

$$\begin{aligned} d\xi_j(t) &= -\kappa \left(\xi_j(t) + {q_j'}^{1/2} \sum_{r\in R} A_{jr} x_r(t)^2 w_r^{-1} \sum_{k\in J} A_{kr} {q_k'}^{1/2} \xi_k(t) \right. \\ &\quad \left. + q_j' \sum_{r\in R} A_{jr} \varepsilon_r^{1/2} x_r(t)^{1/2} dB_r(t) \right), \end{aligned}$$

which is of the form (2.47) where

$$G_{jr} = {q_j'}^{1/2} A_{jr} \varepsilon_r^{1/2} x_r^{1/2},\ j \in J,\ r \in R,$$

and

$$GG^T = {Q'}^{1/2} AEXA^T {Q'}^{1/2},$$

where $E = \mathrm{diag}(\varepsilon_j, j \in J)$.

Other forms of the algorithm may also be obtained where the derivatives remain fixed giving mathematically tractable results.

Finally, we examine the discrete time-lag version of this particular variant of the dual algorithm considered in Section 2.3.4. The system is represented as

$$\mu_j[t+1] = \mu_j[t] + \kappa \left(q_j \Big(\sum_{r\in R} A_{jr} x_r[t - d(j,r)] \Big) - \mu_j[t] \right) \tag{2.63}$$

where the previous definitions still hold. An equilibrium point of the system (2.63), along with equation (2.53), is given by the unique value μ such that the Lyapunov function $\mathcal{V}(x)$ defined by equation (2.42) is maximised. Linearising this system about μ in the usual way, we obtain

$$\xi_j[t+1] = \xi_j[t] - \kappa \left({q_j'}^{1/2} \sum_{r\in R} A_{jr} x_r^2 w_r^{-1} \sum_{k\in J} A_{kr} {q_k'}^{1/2} \xi_k \big[t - d(j,r) - d(k,r) \big] + \xi_j[t] \right). \tag{2.64}$$

In the matrix form equation (2.64) is defined by equations (2.55)–(2.56), where the first row of the matrix M is given by

$$\begin{pmatrix} I-\kappa(L(0)+I) & -\kappa L(1) & -\kappa L(2) & \dots & -\kappa L(D) \end{pmatrix}.$$

We make similar deductions about the stability of the matrix M as in Section 2.3.4.

In this chapter we have shown that a weighted proportionally fair allocation could be achieved by simple rate control algorithms. In particular, we have studied a pair of primal and dual algorithms and their variants. These algorithms provide natural relaxations to the optimisation problems $NETWORK(A, C; w)$ and $DUAL(A, C; w)$ respectively, under properly defined Lyapunov functions. We have identified conditions for stability and the speed of convergence for the system with and without time-lags. The stochastic models measure the magnitudes of fluctuations.

In the remaining chapters, we shall analyse the primal-dual pair of rate control algorithms in greater depth for a class of communication networks.

Chapter 3

Random Networks

The theoretical framework presented in the previous chapter is developed for large-scale communication networks. To this end, we investigate the algorithms for stability and fairness. We select a network which has a flexible structure. Paxson and Floyd [86] argue that there is no such thing as a 'typical' Internet topology because it is constantly changing, and the topology is engineered by a number of different entities. In this chapter, we analyse the rate control algorithms for a random network and contrast the results obtained with those of a network with a definite structure. Since the topology of the communication network having $|J|$ resources and $|R|$ routes is represented by the $|J| \times |R|$ incidence matrix A, a structure of the random network will be determined by A. The matrix A will be therefore random and will satisfy the properties of random matrices. Thus, there is some regularity between the theory of random matrices and communication networks with random structures.

In this chapter, we first review the main results on random matrices developed over the past decades. We then analyse the rate control algorithms introduced in the previous chapter using different approaches. We present each set of the results for one class of algorithm only. Similar deductions can be made from the other class. The eigenvalues of the linearised systems considered satisfy the conditions of symmetric random matrices. We derive conditions for the convergence rates and stability, as well as the magnitude of fluctuations as equilibrium. The chapter ends with a parallel study on a network having a definite structure and the results are compared.

3.1 A Survey of Random Matrices

Random matrices have been studied in many scientific areas since the beginning of this century. Before we go into the analysis of random matrices as applied to communication networks, we present some historical development in this field of study and outline some of the major results. For the interested reader, Mehta [80] and Carmeli [18] provide a good survey on the theory of random matrices.

The origin is from nuclear physics, in the statistical theory of energy levels of heavy atomic nuclei, where it is important to know the distribution of eigenvalues of $n \times n$ matrices $X_n = (X_{ij})$, $1 \leq i, j \leq n$, having independent and identically distributed random variables as elements. Under certain moment conditions, one arrives at the *semicircle law* first proved by Wigner [100] and generalised by Arnold [1, 2]. It states that the density of the eigenvalues of X_n tends, with probability unity to

$$w(x) = \begin{cases} (2/\pi)(1-x^2)^{1/2} & \text{for } |x| \leq 1, \\ 0 & \text{for } |x| > 1. \end{cases}$$

The semicircle law does not in fact fit the observed energy spectra of heavy atomic nuclei for which it was meant, because the spectra seem to have densities which are convex, while the semicircle density is concave. So it is desirable to investigate other classes of Hermitian random matrices.

The recent papers of Arnold [1, 2] and Wegmann [96] treat with rigour many aspects of square symmetric and Hermitian random matrix eigenvalues. Arnold [2] showed that the eigenvalues of a large Hermitian matrix will be approximately distributed according to the semi-circle law, if, roughly speaking, all the X_{ij}'s $(i \geq j)$ are of the same magnitude and are independent of each other.

In a first attempt to analyse randomly connected networks in terms of the spectrum they induce, Grenander and Silverstein [49] introduced a class of neural models corresponding to neural connections of a hypothetical organism at birth. The topology of a network of this type was a symmetric random matrix with dependent elements generated by a controlled probability model. They showed that the resulting linear operator has a spectral measure that converges in probability to a universal one when the size of the net tends to infinity: a *law of large numbers* for the spectra of such operators. In other words, for large networks the local randomness does not influence the spectral measure.

The limiting distribution of the singular values of rectangular random matrices have as many interesting applications to multivariate statistics as the well-known limiting distributions of eigenvalues of square symmetric random matrices have to nuclear physics. Wachter [95] proved the almost sure convergence of the empirical measure of the normalised singular values of increasing rectangular submatrices of an infinite random matrix of independent elements in the more general case where means and variances might vary and only a bound on the moments of order $2+\delta$ for some $\delta > 0$ imposed. With this piece of work, he dispelled the misleading impression of the semicircle law as a canonical limit like the Gaussian distribution.

Several authors [1, 2, 49, 95, 96, 100] have identified limiting spectral distributions for sequences of symmetric real-valued random matrices. Geman [40] investigated the behaviour of the largest eigenvalue, $\lambda_{max}(n)$ $(=\frac{1}{n}|X_{p_n}X'_{p_n}|)$, which cannot be inferred from the limiting spectral distribution, and proved the following result for the norm of random matrices:

$$\lambda_{max}(n) \to (1+\sqrt{y})^2\sigma^2 \text{ almost surely (a.s.),}$$

where the expectations (denoted by $\mathbb{E}$) are $\mathbb{E}X_{11} = 0$, $\mathbb{E}|X_{11}|^n \leq n^{\alpha n}$ for all $n \geq 2$, some α, $\mathbb{E}X_{11}^2 = \sigma^2$ and p_n is a sequence of integers such that

$$\lim_{n\to\infty} \frac{p_n}{n} \to y \text{ for some } 0 < y < \infty.$$

The limiting behaviour of the spectrum of $n^{-1}X_nX'_n$ and of various generalisations have been thoroughly studied and is well understood [95, 100]. The study of the behaviour of $1/\sqrt{n}X_n$ was then motivated by the norm of a random matrix.

Pursuing the work of Geman [40], Bai and Yin [8] considered random matrices which are a power of a square random matrix with identically and independently distributed (i.i.d.) entries and showed that $\limsup_{n\to\infty} \rho_n \leq \sigma$ a.s., with $\rho_n = \lambda_{max}(X_n/\sqrt{n})$ as the spectral radius of $X_n/\sqrt{n}$, $\mathbb{E}X_{11} = 0$, $\mathbb{E}X_{11}^2 = \sigma^2$, the variance of entries of X_n, $\mathbb{E}X_{11}^4 < \infty$, and $\lambda_i(.)$ $i = 1 \dots n$ are the eigenvalues of the corresponding $n \times n$ matrix.

Geman [42] proved the same result under the stronger conditions $\mathbb{E}X_{11} = 0$, $\mathbb{E}X_{11}^2 = \sigma^2$ and $\mathbb{E}X_{11}^n \leq n^{\beta n}$ for all $n \geq 3$ and some $\beta > 0$. He used the fact that the spectral radius of a matrix cannot exceed its Euclidean norm. Hwang [55] verified the real Gaussian case by computer simulations and showed that the limit spectral distribution of $X_n/(\sqrt{n}\sigma)$ distributes uniformly over the unit disc.

The spectra for random matrices of the form $(1/n)X_pX_p^{'}T_p$ are important in many fields. Here $X_p = [X_{ij} : i = 1 \dots p, j \dots n]$ and T_p is a $p \times p$ non-negative definite random matrix independent of X, for $p = 1, 2, \dots$. Many results are available for the special case where $T_p = I_p =$ identity (for example, see Grenander and Silverstein [49], Wachter [95]).

Yin and Krishnaiah [106] showed that the limiting spectral distribution of $(1/n)X_pX_p^{'}T_p$ exists when X_pX_p' is distributed independent of a non-negative definite symmetric random matrix T_p, with i.i.d. and normal entries. A stronger result was obtained in Yin et al. [104] where non-random limits are defined for functions of a central multivariate matrix. This work has been generalised when X_p has *isotropic*[1] elements and their second moments exist (See Bai et al. [9]).

Suppose $\frac{1}{p}tr\ T_p^k \to H_k$ a.s. as $p \to \infty$ for $k = 1, 2, \dots$ and $\sum H_{2k}^{-1/2k} < \infty$. Then Yin [103] proved that the spectral distribution of $A_p = (1/n)X_pX_p^{'}T_p$ tends to a non-

[1] A random vector X_p is said to have the isotropic distribution if the distribution of X_p is the same as PX_p, where P is an $p \times p$ orthogonal matrix.

random limit distribution as $p \to \infty$, $n \to \infty$, but $p/n \to y > 0$ by assuming that X_p has i.i.d. entries and only second moments exist.

Considerable literature is available on the asymptotic distributions of certain functions of eigenvalues of sample covariance matrices. Recently, Bai et al. [7] showed that $\overline{\lim_p} \lambda_{max}(p) = \infty$ a.s. when $\mathbb{E}X_{11}^4 = \infty$. This result verifies the boundedness of $\mathbb{E}X_{11}^4$ to be the weakest condition known to assure the almost sure convergence of $\lambda_{max}(p)$ for a class of sample covariance matrices as compared to the results obtained in [8, 40, 42, 55]. Silverstein [90] has provided necessary and sufficient conditions to establish the convergence in distribution of certain random variables defined by A_p.

It is known when $Var(X_{11}) = 1$, the empirical distribution function of the eigenvalues of A_p converges almost surely for every $x \geq 0$ to a non-random probability distribution function F_y having a density with support on $[(1-\sqrt{y})^2, (1+\sqrt{y})^2]$, and for $y > 1$, F_y places mass $(1 - 1/y)$ at 0 [7, 49, 61, 95, 103]. Moreover, if X_{11} is standardised, $\lambda_{max}(p)$ satisfies $\lambda_{max}(p) \to (1+\sqrt{y})^2$ a.s. as $p \to \infty$, if and only if $\mathbb{E}X_{11}^4 < \infty$ [7, 40].

Yin et al. [105] proved that if $\mathbb{E}X_{11}^4 < \infty$ and $\mathbb{E}X_{11} = 0$, then $\lambda_{max}(p)$ converges a.s. to $(1+\sqrt{y})^2 \mathbb{E}X_{11}^2$ as $n \to \infty$, $p \to \infty$ and $p/n \to y > 0$, while in [7], it was shown that $\limsup_p \lambda_{max}(p) = \infty$ a.s. if $\mathbb{E}X_{11}^4 = \infty$. Silverstein [91] established the weak limiting behaviour of $\lambda_{max}(p)$ and a series of limits for $\lambda_{max}(p)$. Thus, convergence in probability holds only under a slight weakening of the condition needed for almost sure convergence.

Geman has also worked a lot on limit laws for large systems of weakly and randomly coupled equations. In [41], for example, he studied the limiting behaviour of a set of non-linear differential equations with random coefficients and consisting of large *homogeneous* subpopulations and proved a law of large numbers for such a system of differential equations arising in the neural modelling for brainstem respiratory centres. This result was extended by Geman and Hwang [43] to obtain a law of large numbers and a central limit theorem for the solutions of certain algebraic systems, and a weak convergence to a Gaussian process for the solution of a system of differential equations.

In a completely different work, Fang and Krishnaiah [31] investigated the asymptotic joint distributions of certain functions of the eigenvalues of a sample covariance matrix, correlation matrix and canonical correlation matrix in non-normal situations when the population eigenvalues have multiplicities without assuming the underlying distribution to be multivariate normal. Jonnson [61] provide some limit theorems for the eigenvalues of a sample covariance matrix when the dimension of the matrix as well as the sample size tend to infinity. He uses the method of moments to obtain the limit of the cumulative density function of the eigenvalues and shows that the sum of the eigenvalues raised to a certain power is asymptotically normal. Finally, he derived a limit theorem for the *log sum* of the eigenvalues with estimates of expected value and variance. This is a follow up of the work started by Wigner [100] and exploited by a number of researchers.

A major part of this chapter is devoted to the analysis of the rate control algorithms, introduced in Chapter 2, for a random network. The eigenvalues of the linearised systems have been shown to satisfy some of the above results, mainly the boundedness of the largest eigenvalue under the existence of fourth order moments. The eigenvalue analysis is performed for the system of linearised equations for a random network and some interesting features highlighted. We end the chapter with the analysis of the dual algorithm for a network with a specific structure and compare the results. We show that there is some regularity between the choice of the network and the results obtained.

3.2 Analysis of Rate Control Algorithms

The approaches of small-scale simulation often lack the wide mix of traffic and topologies found in real networks. In this section, we shall first describe the choice of the random network for the analysis of the primal and dual rate control algorithms respectively. We also give a main result on the eigenvalues of matrices which supports the analyses. Later, in Section 3.4.3, we describe a regular network which is used to analyse the dual algorithm. The results for the random network and this regular network are then compared for the dual algorithm. These results can also be generalised for the primal algorithms.

We consider a random network where the elements of the $|J| \times |R|$ matrix A are independent Bernoulli random variables, each taking the value 1 with probability ρ, $(0 < \rho < 1)$ and 0 otherwise. The matrix A is therefore a 0–1 matrix generated as

$$A = \{A_{jr} : A_{jr} = Bernoulli(\rho),\ j \in J,\ r \in R\}.$$

Let $w = (w_r = \sum_{j \in J} A_{jr}, r \in R)$ and let $V = (v_j \sum_{s:j \in s} A_{js}, j \in J)$. In this model, the vector w gives the length of the routes, that is, the number of resources through which each route passes, for each source in the network. Similarly the vector V gives the number of routes which passes through each node. The choice of parameters is appropriate for our analysis, since under these assumptions, we get a unique stable solution $x_r = 1$ for each user r and $\mu_j = 1$ for each resource, as desired for both the primal and the dual algorithms respectively. This facilitates the analysis and observations, and can be scaled to suit real networks. Most of the analysis was conducted for a very large random network with $|R| = 1000$, $|J| = 100$ and $\rho = 0.1$, unless otherwise specified.

Before we proceed with the primal-dual pair of algorithms, we give an important result on matrices which is used considerably in the remaining sections of this chapter.

Theorem 3.1 *Let A_n be a real, symmetric matrix of order n such that the diagonal elements are α and the off-diagonal elements are β, where α and β are real numbers. Then the following results hold:*

1. *The eigenvalues of A_n, say, λ_j, $j = 1, 2, \ldots, n$, are*

$$\lambda_j = \begin{cases} \alpha - \beta & \textit{for } j = 1,\ 2,\ \ldots,\ n\text{-}1, \\ \alpha + (n-1)\beta & \textit{for } j = n. \end{cases}$$

2. *The right eigenvector corresponding to the unique eigenvalue $\alpha + (n-1)\beta$ is the n-dimensional unit vector $\boldsymbol{x_n}$ where $\boldsymbol{x_n} = (\underbrace{1, 1, \ldots, 1}_{n})^T$.*

3. *The linearly independent eigenvectors corresponding to the $n-1$ repeated eigenvalues $\alpha - \beta$ are $\boldsymbol{x_i} = (x_{ij},\ j = 1,\ 2,\ \ldots,\ n)^T$, $i = 1,\ 2,\ \ldots,\ n-1$, where*

$$x_{ij} = \begin{cases} -1 & \textit{if } j = 1 \\ 1 & \textit{if } j = n - i + 1 \\ 0 & \textit{otherwise.} \end{cases}$$

Proof.

Proof of result 1. The first result can be easily proved by induction on the order of the matrix A by taking the characteristic polynomial to be

$$\varphi_{k+1}(\lambda) = (\alpha - \lambda - \beta)^k(\alpha - \lambda - k\beta).$$

□

Proof of result 2. The right eigenvector is obtained from

$$[A_n - (\alpha + (n-1)\beta)I]\boldsymbol{x_n} = \mathbf{0},$$

or

$$\begin{bmatrix} -(n-1)\beta & \beta & \dots & \beta \\ \beta & (n-1)\beta & \dots & \vdots \\ \vdots & \ddots & \ddots & \beta \\ \beta & \dots & \beta & -(n-1)\beta \end{bmatrix} \boldsymbol{x_n} = \mathbf{0}.$$

or

$$\begin{bmatrix} -(n-1) & 1 & \dots & 1 \\ 1 & (n-1) & \dots & \vdots \\ \vdots & \ddots & \ddots & 1 \\ 1 & \dots & 1 & -(n-1) \end{bmatrix} \boldsymbol{x_n} = \mathbf{0}.$$

On solving, the above system gives $\boldsymbol{x_n} = (\underbrace{1, 1, \dots, 1}_{n})^T$.

Note that rank of $(A - (\alpha + (n-1)\beta)I)$ is n.

□

Proof of result 3. The corresponding characteristic equation is

$$[A_n - (\alpha - \beta)I]\boldsymbol{x_i} = \mathbf{0}.$$

Since the rank of $[A_n - (\alpha - \beta)I]$ is 1, the linearly independent eigenvectors $\alpha - \beta$, gives $\boldsymbol{x_i} = (x_{ij},\ j = 1, 2, \dots, n)^T$, $i = 1, 2, \dots, n-1$, where

$$x_{ij} = \begin{cases} -1 & \text{if j = 1} \\ 1 & \text{if j = n - i + 1} \\ 0 & \text{otherwise.} \end{cases}$$

■

3.3 The Primal Algorithm

In this section, we study in more detail the effect of the gain parameter κ on the rate of convergence for a class of communication networks. For the random network selected in this dissertation, and described in the previous section, we have the stationary solution of the rates of flow and the shadow prices chosen explicitly to be unity. The choice of parameters for the system leads to $p'_j = 1/v_j$, since all the shadow prices μ_j converge to unity, so that $P' = V^{-1}$. Further, $W = \mathrm{diag}(w_r = \sum_j A_{rj}, r \in R)$, $V = \mathrm{diag}(v_j = \sum_{r:j\in r} A_{jr}, j \in J)$ and $X = I$. With these specifications, the symmetric matrix (2.23) reduces to

$$\Gamma^T \Phi \Gamma = W + A^T V^{-1} A. \tag{3.1}$$

The eigenvalues of the matrix (3.1) were computed for the selected random network with 1000 routes and 100 resources. It was found that $\phi_r \in (1, 25)$ for all $r \in R$ and are distributed as shown in Figure 3.1. The trace of the matrix $\Phi = \mathrm{diag}(\phi_r, r \in R)$

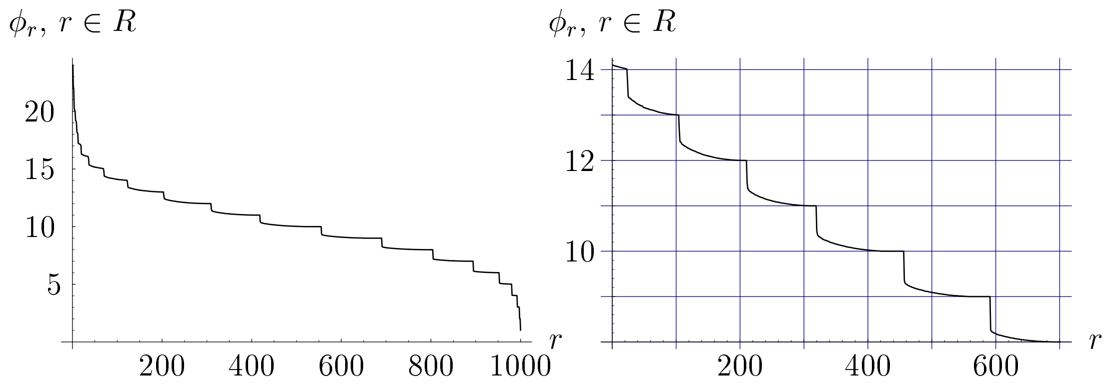

Figure 3.1: Distribution of eigenvalues of the matrix Φ. The eigenvalues corresponding to the large random network are sorted and plotted with the routes in the system. The second plot shows a magnified portion of the first plot. Note the jumps at the integral eigenvalues.

is obtained as 10165. The rate of convergence will be determined by the smallest eigenvalue which, in this case, is 1.0088.

Since the route lengths are at most 22, the question which arises from the plots of the eigenvalue distributions is, whether these eigenvalues are related to the route lengths w_r, $r \in R$. We offer a possible explanation in Section 3.3.1.

Further, from the stochastic analysis of the primal algorithm in Section 2.2.3, it was observed that the covariance matrix increases linearly with the gain parameter κ. As κ increases, the faster convergence to equilibrium is at the cost of a greater spread at steady state.

When there is congestion indication with individual feedback, the matrix F (see Section 2.2.3) is given as

$$F_{r,(j,s)} = \varepsilon_j^{1/2} \mu_j^{1/2} A_{jr} I[r = s], \; r \in R, \tag{3.2}$$

and

$$FF^T = \mathrm{diag}\big(\sum_{j \in r} \mu_j \varepsilon_j, \; r, \; s \in R, \; j \in J \big),$$

where ε_j, $j \in J$, is small.

For the set of parameters for the random network, the shadow prices converge to one at equilibrium. The ε_j can safely be assumed to be unity (since *critical gain*, discussed later in this chapter, is comparatively very small), so that the matrix FF^T reduces to

$$FF^T = \mathrm{diag}\big(\sum_{j \in r} A_{jr}, \; r \in R \big) = W,$$

and

$$[\Gamma F; \Phi]_{rs} = \frac{[\Gamma W \Gamma^T]_{rs}}{\phi_r + \phi_s} \,. \tag{3.3}$$

The smallest and the largest variance, as well as the trace of the matrix $\Sigma = \kappa\Gamma^T[\Gamma F; \Phi]\Gamma$, will describe the spread at equilibrium.

The experiment in this section is performed using a topological matrix A, with $|J| = 10$, $|R| = 100$ and $\rho = 0.5$.

The variance of the matrix Σ lies in the interval $[0.341\kappa, 1.598\kappa]$ and has a trace of 54.9573κ. The result clearly shows the dependency of the variance on the gain parameter: a larger gain implies faster convergence but at the cost of more oscillations at equilibrium. Moreover, at equilibrium, the flow rates will oscillate about the stable point and the magnitude of the fluctuations will vary as per the variance in Σ. A small κ is therefore preferred.

Next, we consider a more general form for the matrix $\Gamma^T\Phi\Gamma$ so that the previous discussion will be a particular case of the matrix considered.
Let

$$\Gamma^T\Phi\Gamma = WX^{-1} + \alpha X^{1/2}A^TP'AX^{1/2} \tag{3.4}$$

where α is a non-negative real constant. The value of α has the effect of changing the slope of P', and hence affect the rate of convergence.

Since $X = I$, equation (3.4) becomes

$$\Gamma^T\Phi\Gamma = W + \alpha A^TV^{-1}A \tag{3.5}$$

which is a real, positive definite, symmetric matrix of order $|R|$. The choice of A ($|J| = 10$, $|R| = 100$, $\rho = 0.5$) ensures that $v_j \neq 0$, $j \in J$, so that V^{-1} exists.

Before we proceed to analyse the system (3.5), we derive a series of simple but useful results about the expectation of the various variables involved.

Lemma 3.1 *Let A be a $|J| \times |R|$ matrix whose elements are i.i.d. Bernoulli random variables with parameter ρ, $0 < \rho < 1$. Further, let $W = \mathrm{diag}(w_r, r \in R)$ and*

$V = \text{diag}(v_j,\, j \in J)$ *where the* w_r*'s are the column sums of* A *and* v_j*'s are the row sums for the same matrix. Let* $W + \alpha A^T V^{-1} A$ *be a real, positive definite, symmetric matrix and* α *be a non-negative real constant. Then the following results hold:*

(i) $\mathbb{E}[w_r] = |J|\rho \;\; \forall\, r \in R$,

(ii) $\mathbb{E}[v_j] = |R|\rho \;\; \forall\, j \in J$,

(iii) $\mathbb{E}[tr(W)] = \mathbb{E}[tr(V)] = |R||J|\rho$,

(iv) $\mathbb{E}[tr(A^T V^{-1} A)] \approx |J|$,

(v) If $\Gamma^T \Phi \Gamma = W + \alpha A^T V^{-1} A$*, where* Γ *is an orthogonal matrix and* Φ *is the diagonal matrix of eigenvalues of* $W + \alpha A^T V^{-1} A$*, then*

$$\mathbb{E}[tr(\Gamma^T \Phi \Gamma)] \approx |J|(|R|\rho + \alpha).$$

Proof. Since $A_{jr} = \{Bernoulli(\rho),\ j \in J,\ r \in R\}$, clearly $\mathbb{E}[A_{jr}] = \rho$. Thus (i) and (ii) follow.

Again, $tr(W) = \sum_{r \in R} w_r$, so that $\mathbb{E}[tr(W)] = |R||J|\rho$. Similarly we have $\mathbb{E}[tr(V)] = |R||J|\rho$, which proves (iii).
Next, the diagonal elements of $A^T V^{-1} A$ are given by

$$[A^T V^{-1} A]_{rr} = \sum_j \frac{A_{jr}^2}{\sum_{r:j \in r} A_{jr}}, \quad r \in R.$$

Proceeding as above, it follows that

$$\mathbb{E}[tr(A^T V^{-1} A)] \approx |R| \mathbb{E}[A^T V^{-1} A]_{rr} = |J|,$$

which is result (iv).

Finally

$$\begin{aligned}\mathbb{E}[tr(\Gamma^T\Phi\Gamma)] &= \mathbb{E}[tr(W)] + \alpha\mathbb{E}[tr(A^TV^{-1}A)] \\ &\approx |J|(|R|\rho + \alpha).\end{aligned}$$

This completes the proof.

■

We can now deduce the main result of the expected eigenvalues of the system (3.5).

Proposition 3.1 *The expected rate of convergence of the generalised system (3.5) is determined by*

$$|J|\rho + (\alpha - \rho)\frac{|J|}{|R|},$$

for any α.

Proof. Taking expectation on equation (3.5) we have

$$\mathbb{E}[W + \alpha A^TV^{-1}A] = \mathbb{E}[W] + \alpha\mathbb{E}[A^TV^{-1}A]. \tag{3.6}$$

But

$$\mathbb{E}[W] = \mathrm{diag}(w_r = \sum_j A_{jr}; r \in R) = \mathrm{diag}(|J|\rho,\ r \in R). \tag{3.7}$$

Consider $\mathbb{E}[A^TV^{-1}A]$. We obtain the expectations for the diagonal and off-diagonal elements separately. In the proof of Lemma 3.1, it was shown that the expectation of the diagonal elements of $A^TV^{-1}A$ is given by

$$\mathbb{E}[A^TV^{-1}A]_{rr} \approx \frac{|J|}{|R|},\ r \in R. \tag{3.8}$$

Similarly, we can show that

$$\mathbb{E}[A^T V^{-1} A]_{kl} \approx \frac{|J|}{|R|}\rho,\ k, l \in R,\ k \neq l. \tag{3.9}$$

Using (3.8) and (3.9) in (3.6), we get

$$\mathbb{E}[W + \alpha A^T V^{-1} A] \approx \frac{|J|}{|R|}\rho \begin{bmatrix} |R| + \alpha/\rho & 1 & \cdots & 1 \\ 1 & |R| + \alpha/\rho & \cdots & 1 \\ \vdots & \ddots & \ddots & \vdots \\ 1 & \cdots & 1 & |R| + \alpha/\rho \end{bmatrix}, \tag{3.10}$$

which is a symmetric matrix with all the diagonal elements are $|J|\rho + \alpha\frac{|J|}{|R|}$ and off-diagonal elements are $\frac{|J|}{|R|}\rho$. Using Theorem 3.1, we get the expected eigenvalues for the system (3.5) as

$$\phi_r = \begin{cases} |J|\rho + (\alpha - \rho)\frac{|J|}{|R|} & \text{for } r = 1 \ldots |R| - 1, \\ 2|J|\rho + (\alpha - \rho)\frac{|J|}{|R|} & \text{for } r = |R|. \end{cases} \tag{3.11}$$

The rate of convergence is decided from

$$\min\left\{|J|\rho + (\alpha - \rho)\frac{|J|}{|R|}, 2|J|\rho + (\alpha - \rho)\frac{|J|}{|R|}\right\},$$

which is clearly $|J|\rho + (\alpha - \rho)\frac{|J|}{|R|}$ for any α.

■

This means that the single largest eigenvalue does not affect the convergence rate. The rate of convergence of a communication network depends on the dimensions of the network and the parameter α.

Remark 3.1 *Since A has dimension $|J| \times |R|$ where $|J| < |R|$, the rank of A, denoted by $rk(A)$, will be less or equal to $|J|$, depending on the value of ρ. Further, since V is* *non-singular, so is P'. Therefore $rk(P') = |J|$. Consequently $rk(A^T P' A) \leq |J|$. This means that $A^T P' A$ can at most have an expected number of $|J|$ non-zero eigenvalues or $|R| - |J|$ null eigenvalues for the particular structure of the network we have chosen for simulation.*

When $\alpha = 1$ we are in fact studying the system (3.1). The eigenvalue analysis for two random matrices of different dimensions are given in Table 3.1. One matrix which is considered has 100 resources and 1000 routes and the topological matrix is generated by $\rho = 0.1$. The second system has 10 nodes and 100 users with $\rho = 0.5$. It can be observed from the table that our analytical results hold in both cases. The results have also been verified for random networks of other dimensions.

Table 3.1: Computational results for the primal algorithm.

	Parameters for Matrix A	
	$\|J\| = 100,\ \|R\| = 1000,$ $\rho = 0.1$	$\|J\| = 10,\ \|R\| = 100,$ $\rho = 0.5$
Range w_r	$[1, 22]$	$[1, 9]$
Range v_j	$[75, 126]$	$[40, 58]$
$\overline{v}$	100.65	50.40
$\overline{w}$	10.065	5.040
$tr(W) = tr(V)$	10065	504
$tr(P')$	1.0050	0.2005
$tr(A^T P' A)$	100	10
$tr(\Phi)$	10165	514
$rk(A^T P' A)$	100	10
ψ_r (rth eigenvalue of $A^T P' A$.)	$[0, 10.950]$	$[0, 5.645]$
ϕ_r	$[1.009, 24.010]$	$[1.015, 12.184]$
$tr(\Sigma)$	—	54.96κ
$\max\{\Sigma_{rr}\}$	—	1.598κ
$\min\{\Sigma_{rr}\}$	—	0.341κ
Expected Theoretical Results		
ϕ_r	$20.09\ \ r = 1000$	$10.05\ \ r = 100$
	$10.09\ \ r = 1 \ldots 999$	$1.09\ \ r = 1 \ldots 99$
$tr(\Phi)$	10100	510

Remark 3.2 *The discrepancies between the theoretical results and the computational results may be allocated to computational accuracy and local randomness. As the matrix becomes large the results have smaller discrepancies. Moreover, a proper choice of the model parameters will give appropriate results for different networks. This conclusion is in accordance with the Law of large numbers for random matrices proved by Grenander et al. [49]. In other words, for large networks the local randomness does not influence the spectral measure.*

3.3.1 Effect of route lengths

In this section, we investigate the relationship between the rate of convergence of the primal algorithm and the length of the routes. The theory developed has been verified experimentally using a network with 100 resources and 1000 routes. Figure 3.2

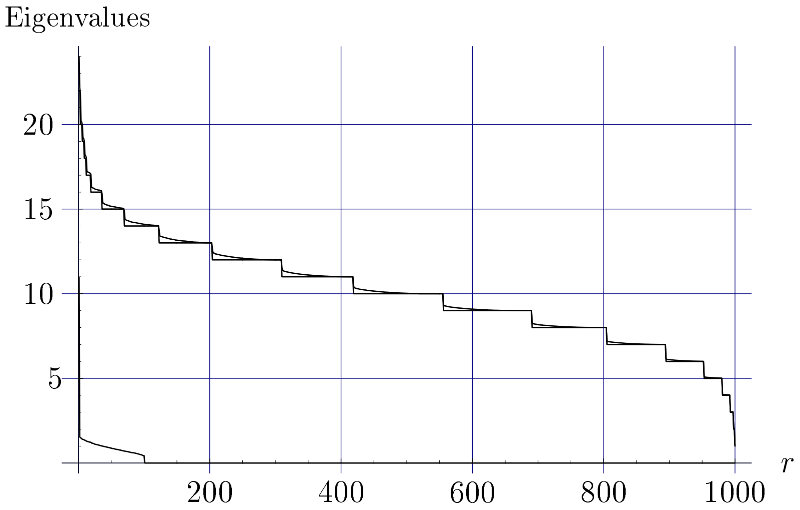

Figure 3.2: Distribution of eigenvalues of the system $\Gamma^T \Phi \Gamma$. The eigenvalues for the different matrices are sorted and plotted against the routes. The lowest cure corresponds to the eigenvalue distribution for the matrix $A^T P' A$; the staircase function corresponds to the eigenvalues of the matrix W; while the third curved superimposed on that for W corresponds to the eigenvalue distribution of the matrix $\Gamma^T \Phi \Gamma$.

shows the distribution of the eigenvalues of the matrices involved in the linearised version of the primal algorithm described on page 54. The behaviour of the route lengths w_r, $r \in R$, is shown as the staircase function with jumps at integral values, representing the change in route lengths. The graph for the eigenvalues of the matrix $\alpha A^T P' A$ for $\alpha = 1$ is the lowest curve in the figure. Since only 10% of its eigenvalues are non-zero[2], a major portion of the curve lies along the x-axis. Moreover, all the non-zero eigenvalues of the same matrix $A^T P' A$ lie in the interval $[0.426, 1.543]$ with the exception of one value which is 10.95. This explains the complete behaviour of this particular function. The distribution of the eigenvalues of the matrix Φ in system (3.5) can be seen as the curve superimposed on the staircase function graph with exponential decrease between any consecutive pairs of integral quantities of the eigenvalues.

Again consider the system (3.5). The matrix of eigenvalues for this system is denoted by $\Phi = \text{diag}(\phi_r, r \in R)$. The characteristic equation of the system under study is

$$|(W + \alpha A^T V^{-1} A) - \phi I| = 0.$$

Since $rk(A^T V^{-1} A) = 100$, and for each resource, v_j, $j \in J$, has between 74 and 127 routes passing through it, the matrix $A^T V^{-1} A$ will have 900 null eigenvalues with the remaining eigenvalues in the interval $[0.426, 1.0523]$. The parameter α is just a scaling factor (in this case one). Moreover, the elements of W (and its eigenvalues) belong to the interval $[1, 22]$. The distribution of the eigenvalues of the matrix $W + \alpha A^T V^{-1} A$ can therefore be explained as follows. The eigenvalues of $A^T V^{-1} A$ will cause small perturbations to the corresponding eigenvalues of W and the joint effect results in the general behaviour in Figure 3.2.

We may assume that $\phi_r \approx \psi_r + w_r$, $r \in R$, where ψ_r is the rth eigenvalue of $\alpha A^T V^{-1} A$. The eigenvalues of the system (3.5) are closely related to the lengths of the various routes of the network under study. This explains the behaviour of the curves in Figure 3.1 at the integer points, as the route lengths are integral values in the interval

[2]As $rk(A^T P' A) = |J|$. See Remark 3.1 on page 57 for an explanation.

[1,22].

The lengths of the routes have therefore a very high contribution to the eigenvalues of the system (3.5) and clearly affect the convergence rate. In other words, *there is a one-to-one correspondence between route lengths and the magnitude of the eigenvalues.* Shorter routes imply lower eigenvalues hence faster convergence rates. Further, the matrix $A^T P' A$ has little effect on the speed of convergence which will increase with P' and hence with the parameter α.

We shall next study the same system for the combined effect of α and κ on the convergence rate. The parameter α mainly measures the spread of oscillations about the stationary point.

Suppose $\phi_r \approx \psi_r + w_r, r \in R$, and consider the system (3.3). Let $\psi_r = \alpha\psi_r',\ r \in R$, where ψ_r' is the rth eigenvalue of $A^T P' A$. Then

$$\begin{aligned} \phi_r &= \psi_r + w_r \\ &= \alpha\psi_r' + w_r,\ r \in R. \end{aligned}$$

Therefore

$$[\Gamma F; \Phi]_{rs} = \frac{[\Gamma W \Gamma^T]_{rs}}{\alpha(\psi_r' + \psi_s') + (w_r + w_s)},\ r, s \in R.$$

Since $\Sigma = \kappa\Gamma^T[\Gamma F; \Phi]\Gamma$, an increase in α will cause a decrease in the covariance matrix Σ, thereby reducing the variability in the system. This will be more so for longer routes. Hence, for a desired control system, there should be a trade-off between w_r, α and κ, assuming the network topology to be fixed.

3.4 The Dual Algorithm

In Chapter 2, we have derived conditions for the stability and rate of convergence for the system of differential equations representing the dual algorithm. We showed that the rate of convergence was proportional to the gain parameter κ and the magnitude of Q' and its convergence to the stable point is determined by the smallest eigenvalue of the matrix $AXW^{-1}XA^T + Q'$. The stochastic versions of this dual algorithm were also studied highlighting the interesting feature that the covariance matrix for the shadow prices increases linearly with the gain parameter κ (Section 2.3.3). The time-lag version of the dual algorithm was also presented in Section 2.3.4.

Next, we analyse the dual algorithm for the random network to determine the relationship between the various parameters in the model in order to achieve optimality and hence maintain stability at equilibrium. We end this section with the analysis of a network with a definite structure and compare the results. Conditions for these specific networks are derived and can be applied to particular classes of communication networks.

3.4.1 Some results for the random network

For the random network selected in this dissertation, we have the stationary solution of the rates of flow and the shadow prices chosen explicitly to be unity. We assume the dual algorithm (2.40)–(2.41) with $q_j = \mu_j(t)v_j$, where $v_j = \sum_{r:j\in r} A_{jr}$. Then $q'_j = v_j$ so that $Q' = \text{diag}(v_j, j \in J)$. The equations representing this system are

$$\frac{d}{dt}\mu_j(t) = \kappa\Big(\sum_{r\in R} A_{jr}x_r(t) - \mu_j(t)v_j\Big) \tag{3.12}$$

and

$$x_r(t) = \frac{w_r}{\sum_{k\in J} A_{kr}\mu_k(t)}. \tag{3.13}$$

Moreover, $W = \text{diag}(w_r = \sum_j A_{rj},\, r \in R)$ and $X = I$. The rate of convergence for the dual algorithm was dealt with in Section 2.3 and is given in the linearised matrix form by equation (2.44). The choice of parameters for the random network gives the particular case of equations (2.44) and (2.45) as

$$\frac{d}{dt}\xi(t) = -\kappa(AW^{-1}A^T + Q')\xi(t) \tag{3.14}$$

with

$$\Theta^T\Psi\Theta = AW^{-1}A^T + Q'. \tag{3.15}$$

Computation of the eigenvalues of the matrix (3.15), which is a square symmetric matrix of order $|J| = 100$, shows that there is a unique eigenvalue with value 206.04 whilst the 99 remaining eigenvalues lie in the range $[80.5, 138.2]$. The trace of Ψ, the matrix of eigenvalues, is then 11069. Therefore the rate of convergence is determined by the minimum of the 99 eigenvalues.

Further, the stochastic analysis of the link algorithm in Section 2.3.3 gave us an expression for the covariance matrix of the system. For shadow prices inferred from fluctuating flow rates, the arbitrary matrix G in equation (2.49) is defined by

$$G_{jr} = \varepsilon_r^{1/2} x_r^{1/2} A_{jr}, \text{ for } j \in J,\ r \in R,$$

with $GG^T = AXEA^T$ where $E = \mathrm{diag}(\varepsilon_r, r \in R)$. Since $X = I$, then GG^T is identically the matrix A, under the assumption that $E = I$, that is, each point of the process contains a workload of unit size. Thus, the covariance matrix obtained from the large random matrix is given by

$$\Sigma = \kappa\Theta^T[\Theta A; \Psi]\Theta \tag{3.16}$$

with

$$[\Theta A; \Psi]_{jk} = \frac{[\Theta A A^T \Theta^T]_{jk}}{\psi_j + \psi_k}. \tag{3.17}$$

The matrix Σ has smallest variance 0.43009κ and largest variance 0.43685κ, for any value of κ. This gives a trace of 43.296κ for Σ. Clearly the variance depends on the gain parameter κ: a larger value of κ implies a faster convergence rate at the cost of more oscillations at equilibrium.

Variability of shadow prices with κ at equilibrium

We study the spread at equilibrium in more detail. It is sufficient to determine the relationship between the mean square deviation of the shadow prices and the gain parameter involved in the dual algorithm. From the simulation results of the deterministic dual case[3], we obtain the *critical value of* κ, denoted by κ^*, as 0.011. This is the value of κ which provides an upper bound for convergence when the spectral radius of the discrete time-lag system is unity. Beyond this point, the system is unstable.

Let the mean of the shadow prices be $\overline{\mu}$ and the corresponding mean square deviation be σ_μ^2. This is defined as the expected value of $(\mu[t]-1)^2$ averaged over all resources. We have

$$\overline{\mu} = \frac{1}{|J|(T_1 - T + 1)} \sum_{t=T}^{T_1} \sum_{j \in J} \mu_j[t]$$

for a time interval of length $(T_1 - T + 1)$ chosen from a simulation run time greater than T_1, and

$$\mathbb{E}\left[(\mu[t]-1)^2\right] = \sigma_\mu^2 = \frac{1}{|J|(T_1 - T + 1)} \sum_{t=T}^{T_1} \sum_{j \in J} (\mu_j[t]-1)^2.$$

The above relations were computed for a very large value of T_1 and $T = T_1/2$ to ensure equilibrium is reached. The output is illustrated in Figure 3.3. A general observation is that $\overline{\mu}$ is slightly above unity, in the range $[1.0003, 1.0022]$, whereas the steady solution for μ is one. This may be explained by the robustness of the shadow prices as compared to the flow rates. For very small values of κ, the relationship

[3]Details about the simulation are given in Chapter 5, Section 5.3.1.

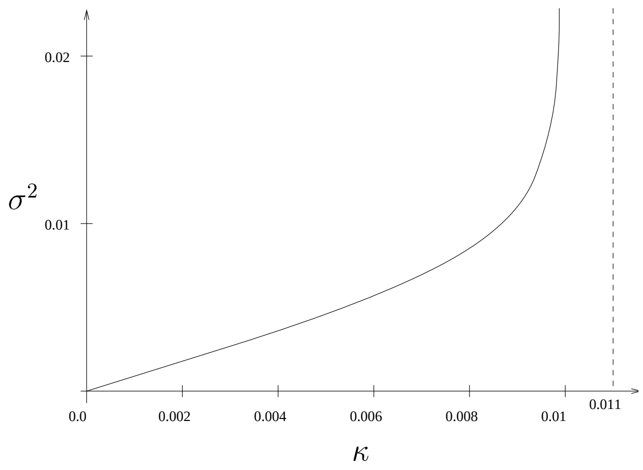

Figure 3.3: Variation of the gain parameter with the mean square deviation of the shadow prices for the resources of the random network. In this case $\kappa^* = 0.11$.

is approximately linear as defined by equation (3.16). In this case, the slope of the linear portion of the curve is approximately 0.5 which compares closely with the mean of the variances obtained in Section 3.4.1. However, as κ increases, the mean square deviation diverges with an asymptote at the critical value of κ where the spectral radii become unity and the deterministic time-lag system becomes unstable.

3.4.2 A variant of the system $AXW^{-1}XA^T + Q'$

In this section we consider a variant of equation (2.45) discussed in Chapter 2. Let

$$\Theta^T \Psi \Theta = AXW^{-1}XA^T + aQ' \tag{3.18}$$

where a is some constant chosen such that the matrix (3.18) remains real, symmetric and positive semi-definite. The parameter a measures the sensitivity of the relationship between flow rates and the shadow prices at the resources. Equation (3.18) reduces to equation (2.45) when $a = 1$, and to equation (3.15) when $X = I$ also. This particular case has been discussed in the previous section. Varying a will change the slope of Q' which in turn will affect the convergence behaviour of the continuous time dual algorithm. We study this effect of varying a on various parameters in the model.

As discussed in Section 3.4.1, for the selected random network, the expression (3.18) reduces to

$$\Theta^T \Psi \Theta = AW^{-1}A^T + aQ' \tag{3.19}$$

which is of order $|J|$ and where $Q' = \mathrm{diag}(v_j, j \in J)$.

Consider the expectation of the matrix (3.19).

$$\mathbb{E}[AW^{-1}A^T + aQ'] = \mathbb{E}[AW^{-1}A^T] + a\mathbb{E}[Q'].$$

Also,

$$\mathbb{E}[Q'] = \mathrm{diag}(\mathbb{E}[v_j], j \in J) = |R|\rho$$

and

$$\mathbb{E}[AW^{-1}A^T]_{jj} \approx \frac{\sum_r A_{jr}A_{jr}}{\sum_j A_{jr}} = \frac{|R|}{|J|}.$$

Again

$$\mathbb{E}[AW^{-1}A^T]_{jk} \approx \frac{\sum_r A_{jr}A_{kr}}{\sum_j A_{jr}} = \frac{|R|}{|J|}\rho.$$

Therefore

$$\mathbb{E}[AW^{-1}A^T + aI] \approx \frac{|R|}{|J|}\begin{bmatrix} 1+a|J|\rho & \rho & \cdots & \rho \\ \rho & 1+a|J|\rho & \cdots & \rho \\ \vdots & & \ddots & \vdots \\ \rho & \cdots & \rho & 1+a|J|\rho \end{bmatrix} \tag{3.20}$$

where $0 < \rho \leq 1$ and a is a constant to be determined.

The rate of convergence of the continuous time dual algorithm is determined by the smallest eigenvalue of the matrix (3.19). We observe that the matrix (3.20) is of the particular form in which the diagonal elements are $\dfrac{|R|}{|J|} + a|R|\rho$ and off-diagonal elements $\dfrac{|R|}{|J|}\rho$. Using Theorem 3.1 in Section 3.2, we deduce that the expected eigenvalues of the matrix (3.19), say, $\psi_1, \psi_2, \ldots, \psi_{|J|}$, are as follows:

$$\begin{aligned}\psi_i &= \left(\frac{|R|}{|J|} + a|R|\rho\right) - \frac{|R|}{|J|}\rho \\ &= \frac{|R|}{|J|}(1-\rho) + a|R|\rho \ \text{ for } \ i = 1, 2, \ldots, |J|-1,\end{aligned} \tag{3.21}$$

and

$$\begin{aligned}\psi_{|J|} &= \left(\frac{|R|}{|J|} + a|R|\rho\right) + (|J|-1)\frac{|R|}{|J|}\rho \\ &= \frac{|R|}{|J|}(1-\rho) + (a+1)|R|\rho.\end{aligned} \tag{3.22}$$

This means that we have a unique eigenvalue as given by (3.22) and $|J|-1$ eigenvalues as given by (3.21). Further, the trace of the matrix of expected eigenvalues becomes $|R|(1 + a|J|\rho)$.

The rate of convergence is established from

$$\min\left\{\frac{|R|}{|J|}(1-\rho) + a|R|\rho, \frac{|R|}{|J|}(1-\rho) + (a+1)|R|\rho\right\}.$$

which is simply $\dfrac{|R|}{|J|}(1-\rho) + a|R|\rho$ for any value of a. This corresponds to the $|J|-1$ eigenvalues. Thus, the unique eigenvalue (3.22) does not affect the convergence rate. In fact, the rate of convergence increases with a, and hence with the slope of Q'.

For the random network under study, $|R| = 1000$, $|J| = 100$, and $\rho = 0.1$, so that the minimum of the eigenvalues is expected to be $9 + 100a$.

For the particular case $a = 1$, we expect the unique eigenvalue to be 209, approximately, and the remaining $|J|-1$ eigenvalues to have a mean of 109with the matrix of eigenvalues having a trace of 11000. Thus the smallest of $|J|-1$ eigenvalues determines the convergence rate. The above calculations compare very well with the results obtained in Section 3.4.1, with the 99 eigenvalues having a mean of approximately 109. The discrepancies are due to computational error and the fact that the theoretical results are based on the expectation of matrix (3.19).

Again from Theorem 3.1 in Section 3.2, the eigenvector corresponding to the unique eigenvalue (3.22) is given by the unit vector $\boldsymbol{x}_{|\boldsymbol{J}|} = (\underbrace{1, 1, \ldots, 1}_{|J|})^T$ whereas for the remaining $|J| - 1$ eigenvalues, the corresponding eigenvectors are given by $\boldsymbol{x_i} = (x_{ij};\, j = 1, 2, \ldots, |J| - 1)^T$ for $i = 1, 2, \ldots, |J| - 1$, where

$$x_{ij} = \begin{cases} -1 & \text{if } j = 1 \\ 1 & \text{if } j = |J| - i + 1 \\ 0 & \text{otherwise.} \end{cases}$$

Thus the $|J| - 1$ eigenvectors are rotations of each other.

The characteristic equation of the matrix (3.19) suggests that the larger the value of a, the faster the rate of convergence to the stationary solution.

Relationship between a and κ

We have seen how a affects the rate of convergence. We next study how a and κ are related. Recall that by critical value of κ we mean that value of κ which provides an upper bound for convergence when the spectral radius of the discrete time-lag system is unity. The critical values of κ for the dual model with time-lags are obtained by simulating the discrete version of the deterministic dual algorithm with time-lags for various values of a and κ (see Chapter 5 for the simulation results). The algorithm used here is the discrete form of equation (3.12) and is given by

$$\mu_j[t+1] = \mu_j[t] + \kappa \left(\sum_{r \in R} A_{jr} x_r[t - d(j,r)] - \left(a(\mu_j[t] - 1) + 1\right) v_j \right), \qquad (3.23)$$

with $x_r[t] = x_r[0]$ if $t < 0$.

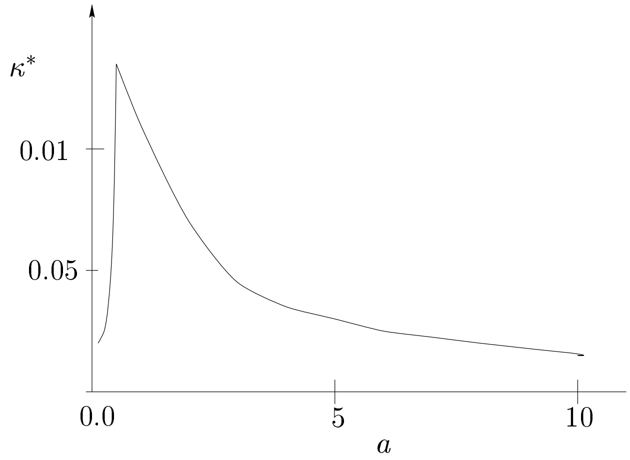

Figure 3.4: Critical gain parameter against a for a random network.

A graph of the critical values of κ with a is given in Figure 3.4. We observe that the curve reaches a maximum at $\kappa = 0.0136$ for $a = 0.5$, and decreases exponentially for other values. As $a \to 0$ (when Q is constant), equation (3.23) reduces to

$$\mu_j[t+1] = \mu_j[t] + \kappa \left(\sum_{r:j \in r} x_r[t - d(j,r)] - v_j \right).$$

The increase in $\mu_j[t]$ then depends on the second term of the above expression which explains the behaviour near the gain parameter axis. In fact, in practice, for an optimal system it would be convenient to use a higher value of κ, to give a faster rate of convergence. In the next chapter, we explore the relationship between convergence times and critical gain parameters.

Effect of κ on the spread at equilibrium as a varies

The previous discussion encourages us to investigate how the variation of the mean square deviation (defined in Section 3.4.1) with κ is affected by a. The same formulae are used, being the particular case when $a = 1$. Figure 3.5 shows this situation for $a = 0.5, 1, 2$ and 5.

As a increases (for $a \geq 1$), the mean square deviation grows exponentially faster for

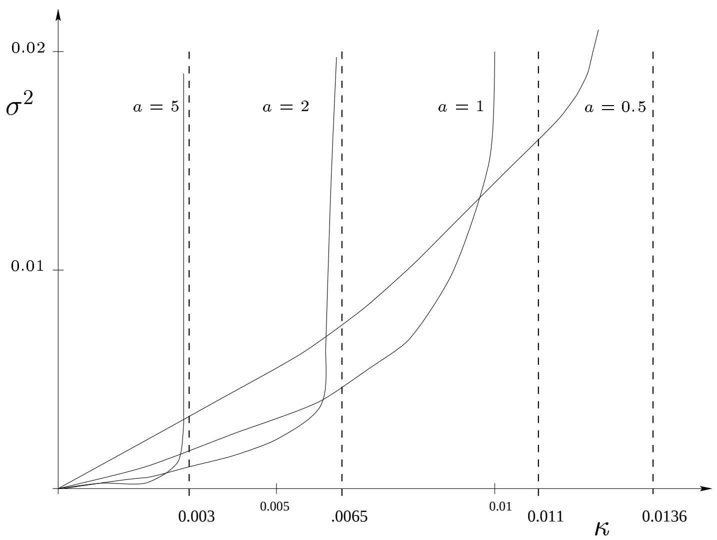

Figure 3.5: Relationship between σ_μ^2 and κ as a varies for a random network. Corresponding to $a = 0.5, 1, 2, 5$, the critical gains are $0.0136, 0.011, 0.0065, 0.003$ respectively.

smaller values of κ. The reason is that as a increases, the value of κ^* decreases, as shown in Figure 3.4. For very small values of κ, an increase in a causes the mean square deviation to increase linearly with κ at a smaller rate, as expected from equation (2.49). In other words, the variability decreases for smaller value of κ and higher values of a. The critical values of κ for each a become asymptotes for the mean square deviation. These represent the situation when the spectral radius approaches unity so that the deterministic time-lag system becomes unstable. For $a = 0.5$, the linear increase of σ_μ^2 can be clearly seen for small κ. Using Figure 3.4, it can be deduced that a further decrease in a will move the asymptote towards the y-axis in Figure 3.5. From the derivation of equation (2.49), it can be easily shown that the matrix (3.18) will cause the matrix Σ to vary inversely with a. The slope of the corresponding curve is then expected to become higher, so that the system is no longer having a smaller variability. This situation is undesirable. Thus, for the chosen network it is convenient to select a to be at least 0.5.

The computations for the matrix (3.19), corresponding to a large random network

Table 3.2: Summary of results for the dual algorithm in a random network.

Numerical results				
a	0.5	1	2	5
$\psi_{\|J\|}$	153.69	206.04	312.53	654.60
$\psi_1 \dots \psi_{\|J\|-1}$	$[43.25, 75.23]$	$[80.50, 138.17]$	$[154.70, 264.31]$	$[376.89, 634.51]$
$tr(\Psi)$	6034.5	11069	21138	51345
$tr(\Sigma)$	78.21κ	43.30κ	23.05κ	9.65κ
$\max\{\Sigma_{jj}\}$	0.793κ	0.437κ	0.232κ	0.097κ
$\min\{\Sigma_{jj}\}$	0.773κ	0.430κ	0.230κ	0.096κ
Slope of curve	0.96κ	0.54κ	0.27κ	0.12κ
Expected theoretical results				
$\psi_{\|J\|}$	159	209	309	609
$\psi_1 \dots \psi_{\|J\|-1}$	59	109	209	509
$tr(\Psi)$	6000	11000	21000	51000

is summarised in Table 3.2. It can be seen that the theoretical eigenvalues given by equations (3.21)–(3.22) are comparable to the numerical eigenvalues for this particular network, the discrepancies being allocated mainly to computational error and local randomness [49]. As a increases, the difference between the two largest eigenvalues decreases, that is, there is more spread of the eigenvalues. Convergence is determined by the smallest eigenvalue while the large eigenvalues are responsible for the oscillations. As a increases, although convergence occurs for all the nodes, there is more jittering. The size of these oscillations increases with a. Further, the slopes of the linear portions of the curve in Figure 3.5 computed for very small values of κ approximates the variance in matrix (2.49). Hence, the numerical results for the large random network strongly supports the theoretical analysis of the dual algorithm.

3.4.3 A regular network

The analysis of the random network has shown many features, some of which can be clearly explained while others are still unclear. To study some of the facets which might have been overlooked in the random network case and to explain some phenomena observed therein, it is sensible to study some networks having definite structures. In this section we consider an example of such a network, which we call a *regular* *network.*

Consider a network having $|J|$ resources, and let C_j be the finite capacity of resource j, for $j \in J$. We assume the flow to be unidirectional as shown in Figure 3.6. Define

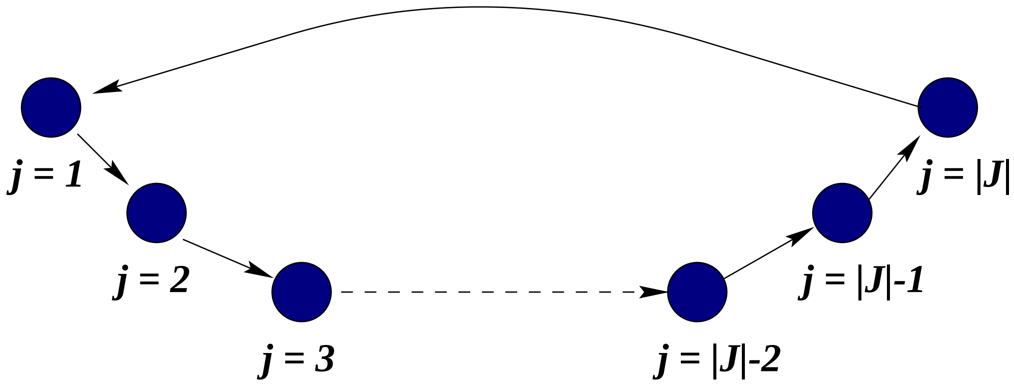

Figure 3.6: A regular network.

a route r in the network as an ordered pair of resources $\{i, j\}$, $i, j \in J$, such that the flow along r is in the direction i to j. Then the number of possible routes, $|R|$, in the network is $|J|^2$. The structure of this network is then defined as follows:

$$A_{k,\{i,j\}} = A_{k,r} = \begin{cases} 1 & \text{if } i \leq j \text{ and } \quad i \leq k \leq j \\ 1 & \text{if } i > j \text{ and } \quad \text{either } j \geq k \text{ or } k \geq i \\ 0 & \text{otherwise.} \end{cases} \tag{3.24}$$

This defines the 0–1 incidence matrix $A = (A_{jr}, j \in J, r \in R)$ for this regular network which is of order $|J| \times |R|$ consisting of $|J|^3$ number of elements. Let $w_r = \sum_k A_{k,r}$, $r \in R$, be the charge to be paid by user r to the network, that is, this charge depends on the route length. Thus, the vector $w = (w_r, r \in R)$ gives all the possible routes in the network. Similarly, $v_k = \sum_r A_{k,r}$, for $j \in J$, gives the number of nodes passing through a particular resource of the network. Let $V = (v_j, j \in J)$. For this network, v_j is fixed for all $j \in J$ and is given by $\displaystyle\sum_{l=1}^{|J|} l = \frac{|J|(1 + |J|)}{2}$. Moreover, for the simulation of the discrete time-lag models we include a delay matrix d in the algorithms used for this network. The definitions we give to the elements of d are given by

$$d_{k,\{i,j\}} = d_{k,r} = \begin{cases} k-i & \text{if } i \leq j \text{ and } \; i \leq k \leq j \\ J+k-i & \text{if } i > j \text{ and } \; \text{either } j \geq k \text{ or } k \geq i \\ 0 & \text{otherwise.} \end{cases} \tag{3.25}$$

Under the above definitions, we get a unique stationary solution $x_r = 1$, $r \in R$, and $\mu_j = 1$, $j \in J$, for the primal and the dual algorithms respectively.

Analysis of the dual algorithm for the regular network

A similar analysis as for the case of the random network (Section 3.4) is performed on the regular network. With the choice of parameters for the regular network, we have $X = I$, $E = I$, and

$$Q' = \text{diag}(v_j, j \in J) = \left(\sum_{l=1}^{|J|} l \right) I_{|J|}$$

since v_j is fixed for all j belonging to J. This reduces the matrix $AXW^{-1}XA^T + Q'$ to $AW^{-1}A^T + \gamma I$ which is of the order $|J|$ and where $\gamma = \sum_{l=1}^{|J|} l$.

Equation (3.18) then reduces to

$$\Theta^T \Psi \Theta = AW^{-1}A^T + a\gamma I_{|J|}. \tag{3.26}$$

Its characteristic equation is therefore

$$|AW^{-1}A^T + a\gamma I_{|J|} - \psi I| = 0$$

or

$$|AW^{-1}A^T + (a\gamma - \psi)I| = 0.$$

Thus, $(a\gamma - \psi)_j$ are the eigenvalues of $AW^{-1}A^T$. Let the eigenvalues of $AW^{-1}A^T$ be θ_j, $j \in J$. Then $\psi_j = \theta_j + a\gamma$. Since A is fixed, W^{-1} is fixed and $AW^{-1}A^T$ is a scalar. This implies that θ_j and γ are fixed, so that the corresponding eigenvalues increase linearly in a. Hence the larger the value of a, the faster the rate of convergence.

Since $\Psi = \mathrm{diag}(\psi_j, j \in J)$, we have

$$\begin{aligned} tr(\Psi) &= \sum_{j\in|J|}(\theta_j + a\gamma) \\ &= |J|a\gamma + \sum_{j\in|J|}\theta_j. \end{aligned}$$

When $|J| = 10$, we have $\gamma = 55$ so that $\mathrm{tr}(\Psi) = 550a + \sum_j \theta_j$.

Further, the ith diagonal element of $AW^{-1}A^T$ is given by

$$(AW^{-1}A^T)_{ii} = \sum_{k=1}^{R}\left(\frac{A_{ik}^2}{\sum_{j\in k} A_{jk}}\right) = |J|.$$

Therefore, $\mathrm{tr}[AW^{-1}A^T] = |J|^2$ and

$$\begin{aligned} tr(\Psi) &= |J|(a\gamma + |J|) \\ &= 550a + 100. \end{aligned}$$

The above simple analysis shows that there will be a unique large eigenvalue $\psi_{|J|} = \theta_{|J|} + 55a$. The other eigenvalues apart from the minimum root which is $\psi_1 = \theta_1 + 55a$, occur in pairs in a smaller range from ψ_2 to $\psi_{|J|-1}$ such that $\max(\psi_2, \ldots \psi_{|J|-1}) < \psi_{|J|}$. This result follows from Theorem 3.1.

The unit eigenvector corresponding to the unique large eigenvalue $\psi_{|J|}$ is $x_{|J|} = \underbrace{(1, 1, \ldots, 1)}_{|J|}$. The unit eigenvector corresponding to the smallest eigenvalue $\psi_1 = \theta_1 + 55a$ is $\underbrace{(-1, 1, -1, \ldots, 1, -1)}_{|J|}$, a vector of 1 and -1 alternatively. The remaining eigenvectors correspond to pairs of equal eigenvalues occurring linearly independently. For this regular network, the eigenvalues of $AW^{-1}A^T$ are: $\theta_{10} = 55$; $\theta_9 = \theta_8 = 12.6774$; $\theta_7 = \theta_6 = 4.4107$; $\theta_5 = \theta_4 = 2.5698$; $\theta_3 = \theta_2 = 1.9483$; $\theta_1 = 1.7873$. Observe that the sum of the eigenvalues is 100 as expected. The minimum eigenvalue of the system is $\psi_1 = 1.7873 + 55a$ and varies with a, so that the rate of convergence varies with a. Using equation (3.18) for this regular network, the covariance matrix Σ is obtained as

$$\Sigma = \kappa\Theta^T[\Theta G; \Psi]\Theta$$

where $\Psi = \text{diag}(\psi_j, j \in J)$ where $\psi_j = \theta_j + a\gamma$, Θ is an orthogonal matrix, $\Theta\Theta^T = I$ and

$$\begin{aligned} [\Theta G; \Psi]_{jk} &= \frac{[\Theta AA^T\Theta^T]_{jk}}{\psi_j + \psi_k} \\ &= \frac{[\Theta AA^T\Theta^T]_{jk}}{(\theta_j + \theta_k) + 2a\gamma} \end{aligned} \tag{3.27}$$

so that an increase in a and/or γ will cause Σ to decrease.

We have seen how a affects the rate of convergence. Next, we study how a and κ are related. We use the same algorithm (3.23) in Section 3.4.2. The plot for the critical values of κ with a is given in Figure 3.7. We observe that the curve reaches a maximum at $\kappa = 0.03$ for $a = 0.2$, and decreases exponentially for other values. This behaviour agrees with that in Figure 3.4 and discussed in Section 3.4.2. In fact, for

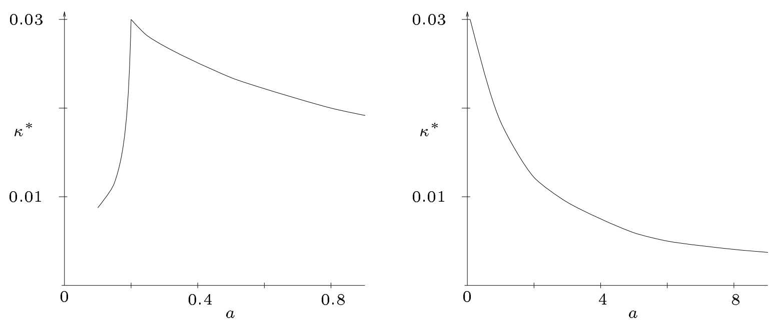

Figure 3.7: Variation of κ^* with a for a regular network.

an optimal system it would be convenient to use a higher value of κ, to give a faster rate of convergence.

As for the random network, we also study the variation of the mean square deviation for the shadow prices with κ as a changes. We observe a similar pattern as in Figure 3.5, on different scales. For this regular network, the critical gains for $a = 0.2, 0.5, 1, 2$ and 5 are $\kappa^* = 0.03, 0.024, 0.0183, 0.0121$ and 0.0052, respectively. Therefore, it is convenient to select a to be at least 0.2 for the regular network.

The computations for the matrix (3.26), corresponding to the regular network is summarised in the Table 3.3. As a increases, the variance decreases together with the slope of the linear portion of the curve. This implies that there is less variability for larger values of a. Furthermore, the stability interval becomes smaller for larger values of a. The slope of the curves for small values of κ are approximately twice the

Table 3.3: Summary of results for the dual algorithm in the regular network.

a	0.5	1	2	5
$\psi_{\|J\|}$	82.5	110	165	330
$\psi_1 \dots \psi_{\|J\|-1}$	$[29.29, 40.18]$	$[56.79, 67.68]$	$[111.79, 122.68]$	$[276.79, 287.68]$
$tr(\Psi)$	375	650	1200	2850
$tr(\Sigma)$	4.700κ	3.040κ	1.870κ	0.880κ
$\max\{\Sigma_{jj}\}$	0.502κ	0.309κ	0.196κ	0.089κ
$\min\{\Sigma_{jj}\}$	0.449κ	0.300κ	0.181κ	0.086κ
Slope of curve	0.810κ	0.590κ	0.350κ	0.160κ

corresponding variances.

We end this chapter with a brief discussion on the main results. The linearised versions of the primal and dual algorithms, presented in Chapter 2, were analysed. Since both algorithms can be analysed in the same way, we have presented our results for one class of algorithms. The other class can be dealt with using the same approach, with similar conclusions.

The eigenvalue analysis of general forms of the linearised primal-dual pair of algorithms have revealed that the rate of convergence of the flow rates and shadow prices are determined by $|J|\rho + (\alpha - \rho)\frac{|J|}{|R|}$, and $\frac{|R|}{|J|}(1-\rho) + a|R|\rho$, for any values of a and α, respectively. This means that the network topology has an important rôle to play in the stability of a system at equilibrium. We have also shown that shorter routes have faster convergence.

The variability of the shadow prices at equilibrium have been shown to depend on the gain parameters as well as on the sensitivity factor a (or α). A large value of a (or α) gives a fixed variance, which implies that the fluctuations are of the same magnitude at equilibrium. This situation gives rise to the phenomenon of *phase-locking* [37]. On the other hand, a smaller a (or α) implies a larger κ^* and hence more variance about the stable point. This in turn gives rise to a faster convergence at the cost of more spread.

The dual algorithm was simulated for both a random network and a regular network. The results are striking and show that they hold equivalently for both types of networks. In particular, for the deterministic regular network, the results about the eigenvalues of the general linearised system are clearly dependent on the sensitivity factor a, and the number of resources in the network. Therefore, there is definitely some regularity between the results obtained for the two types of networks. It remains to verify the generalisation of the results for any communication network.

The results allow us to quantify the various relationships and trade-offs between the rate of convergence of the algorithms, the magnitude of fluctuations about the stable point, and the stability of the system.

The analysis of the primal rate control algorithms can equally be performed for the dual algorithms and vice-versa. Further, for the models studied in this chapter, we assumed the gain parameters to be fixed and no explicit results could be made when κ is different for each user. It would be interesting to know if stability of the time-lag systems is better when κ depends on the lengths of the routes. In the next chapter we investigate a model in which the gain parameter is route dependent.

Chapter 4

A Single Bottleneck Resource

In Chapter 2, we established stability conditions for a general communication network using a Lyapunov function approach. We observed that, not only does the speed of convergence increase with the gain parameter at the cost of more spread at equilibrium, but it is also affected by the smallest eigenvalue of the real, symmetric matrix representing the linearised system. These aspects were studied in greater depth in Chapter 3. In the presence of time-lags, however, increasing the gain parameter may compromise the stability of the system. It was noted that the stability of the primal and dual time-lag algorithms are determined by the spectral radii of the matrices (2.39) and (2.56), respectively. Moreover, the situation when the gain parameter was not fixed did not reveal any explicit results from the analysis in Chapter 3.

The aim of this chapter is to investigate the stability of a network with a single bottleneck resource and two users when the gain parameters as well as the round-trip times for each connection vary. The model in Chapter 2 ignored round-trip times within the network, in order to focus on the dynamical behaviour over longer time periods. In the present context, we consider stability on shorter time-scales, comparable to the round-trip times in the network.

In the current Internet, the rate at which a source sends packets is controlled by TCP, the Transmission Control Protocol of the Internet [57], implemented as software on the computers that are the source and destination of the data. The general approach [24] is as follows. When a resource within the network becomes overloaded, one or more packets are lost; loss of a packet is taken as an indication of congestion, the destination informs the source, and the source slows down. The TCP then gradually increases its sending rate until it again receives an indication of congestion. This cycle of increase and decrease assists to discover and utilise whatever bandwidth is available, and to share it between flows. The steady-state behaviour of Jacobson's transmission control protocol has been analysed extensively. At equilibrium, the rate allocated to a user is approximately given by $x \approx \dfrac{c}{\tau\sqrt{p}}$, where τ is the round-trip time of the connection, p is the packet loss probability over the connection, and c is a known constant. For a detailed discussion on this relationship, the reader is referred to [36, 45].

Presently, there is an incentive to modify the TCP algorithm so that it strives more aggressively for a larger share of available bandwidth, or even to avoid using any form of congestion control. A body of work is now emerging that takes a radically different approach to differentiated services (see, for example, [28, 45, 69]). Its premise is that a simple packet network may be able to support an arbitrarily differentiated set of services by conveying information on congestion from the network to intelligent end-nodes, which themselves determine what should be their demands on the packet network. If the resource implications of their actions can be made known to the end-nodes, then the latter are best placed to determine what should be their demands upon the resources of the network. The optimal allocation of resources depends upon the utilities the various users attach to their several flows through the network, as well as upon the properties of the resources within the network.

The approach in [28] has demonstrated the possibility of a *charge-aware* TCP, whereby the congestion parameters are alterable by end-nodes in an incentive-compatible manner. Just as the end-nodes may choose the willingness-to-pay w to suit their needs, they may have different requirements for the gain parameter. In this chapter, we investigate how the end-users may choose these parameters in order to maintain the stability of the system. In particular, we focus on gain parameters and round-trip times.

Packet marking strategies are familiar to flow control specialists, and are an active research area, see for example [25, 33, 38, 44, 45]. The simple network mechanism, of setting just a single bit to mark some packets [33], may be enough to implement a *smart mechanism* [76] for efficient allocation of resources; since, as end-systems see more packets per round-trip time, any randomness associated with whether or not a particular packet is marked becomes less relevant than the proportion of packets marked. The marking mechanism proposed by Gibbens and Kelly [45] will generally produce bursts of marked packets, separated by longer periods when no packets are marked. An end-node's response to such fluctuations will be smoothed by, for example, any gain parameter κ used by the end-node. It seems advantageous to conduct such smoothing at end-nodes rather than at resources; each end-node will have the ability and the incentive to relate its own smoothing behaviour to its own feedback delay. The marks provide a rational basis for a usage-related charging scheme.

An important feature of Jacobson's algorithm [57] is that it is *self-clocking*. Self-clocking may be chosen as a sensible strategy for routes with uncertain round-trip times; it certainly simplifies the choice of the gain parameter κ, which might otherwise be expected to depend upon an estimate or bound on the round-trip time of a connection. We investigate this aspect further in this chapter.

The recent work of Key et al. [70] introduces a marking mechanism implemented on the primal algorithm in Section 2.2 for a single resource. They argue that in the algorithm, there can be an implicit dependence through κ_r, which might be related to the reciprocal of the round-trip time. Replacing the willingness-to-pay w_r by $\frac{w_r}{\tau_r}$, where τ_r is the round-trip time for user r, would make the steady state throughput dependent on (the reciprocal of) the round-trip time. They raise another question about the choice of κ and $x(0)$, the initial rate, and intuitively suggest that a 'good' choice will be a function of the network topology (example, round-trip time). In our present work, we assume that the gain parameters are adaptive, in the sense that the users (assumed to be adaptive) do make an effort to choose the gain parameters in such a way that they have a proportionally fair share of the bandwidth while maintaining stability at equilibrium.

Important recent papers of Bolot and Shankar [15], Fendick et al. [32], and Bonomi et al. [16] have analysed the stability of networks with a single bottleneck resource, where congestion is signalled by the build-up of a queue at the bottleneck's buffer, and where propagation delays are significant. They have established local instability of the equilibrium point whatever gain parameters are used. The framework we adopt in this chapter is the new concept of marking proposed by Gibbens and Kelly [45], and implemented in a variant of the primal algorithm discussed in Section 2.2.

This chapter is organised as follows. In the present context, we shall consider the impact of round-trip times on stability in a very simple case. In the next section, we present the single bottleneck resource network as a system of delay-differential equations. The model is a particular case of the system described in Chapter 2. However, the model in Chapter 2 ignored round-trip times within the network, in order to focus on the dynamical behaviour over longer time periods. The theoretical analysis of the deterministic version of the system is performed in Section 4.2. This is followed by the analysis of the discrete time-lag version of the model in Section 4.3. We linearise the system by assuming small perturbations about the stable point and obtain stability conditions for the time-lag model. Section 4.4 reviews some stability bounds while in Section 4.5, we present some numerical results based on the derivations in the previous sections. Interesting results are obtained concerning the choice of the gain parameters by the end-nodes to maintain stability while being allocated an optimal fair share of the bandwidth. The convergence times of the discrete time-lag system are studied, for different scenarios as the round-trip time changes, in Section 4.6. We claim that a proper choice of the gain parameters does affect the convergence times, and hence stability.

4.1 The Network Model

Consider a system with a single bottleneck resource and R end-nodes. Let the rate of flow for user r, $r \in R$, at any time point t be $x_r(t)$ when the route r has a round-trip time τ_r. Suppose that the arrival rates of the packets follow a Poisson process and feedbacks indicating congestion are sent to the users. Assume that the resource has capacity to cope with N equally sized packets. Congestion at the link is measured by queue length. In an environment with packets of different sizes, the model can be slightly modified to be compatible with our framework [45]. Each packet is marked with a probability which increases with the average queue length. By appropriately marking packets at overloaded resources, end-nodes are provided with the necessary information to make efficient use of the network. When the mean queue length is less than a defined threshold N, no packets are marked. In case the number of packets arriving in a time unit exceeds N, a mark is placed on *each* of these packets. (We shall not distinguish between packets which are lost and those which are merely marked: we treat them all as marked packets.) The rate of sending packets increases linearly at a rate proportional to the willingness-to-pay w_r for the user r until a feedback signal is received indicating congestion. At that point, the source decreases the rate of packet sending multiplicatively at a rate proportional to the rate at which the packets are being marked at the bottleneck resource.

The differential equations representing the single-resource system, taking time-lags into account, are given by

$$\frac{d}{dt}x_r(t) = \kappa_r\bigg(w_r - x_r(t-\tau_r)\mu(t-\tau_r)\bigg), \quad r \in R, \tag{4.1}$$

where

$$\mu(t) = p\bigg(\sum_{s\in R} x_s(t)\bigg). \tag{4.2}$$

Equations (4.1)–(4.2) represent a TCP-like algorithm for Internet traffic [57, 67], where κ_r is the gain parameter for user r, τ_r is the round-trip time of connection r, and $\mu(t)$ is the shadow price for the bottleneck resource in the system. The system of equations may be motivated in as follows: suppose that $p(y)$ is the price charged by the resource, per unit flow through that single resource, when the load through that node is y. Then, the network attempts to equalise the aggregate cost of this flow, $x_r(t)\mu(t)$, with a target value w_r, by adjusting the flow $x_r(t)$ on route r according to equations (4.1)–(4.2), [26, 57, 68].

The function $p(y)$ is assumed to be non-negative and continuous increasing in y. This function can take different forms depending on how congestion is measured [45]. We interpret $p(y)$ as the shadow price of the resource: it is the marginal increment in expected cost at the resource for a marginal increment in load. Then, the expected charge per unit time to flow r, $x_r p(y)$, is precisely the fair charge to this flow under the model in Chapter 2, where x_r is the optimal flow rate at equilibrium for user r. Motivations for different marking strategies, and hence for different functions $p(.)$, are many and varied (the interested reader is referred to [101]). In this chapter, we assume that $p(y) = y^N$. This is the probability of marking all packets exceeding the threshold N. This specific function helps us understand the impact of threshold levels and traffic variability at the packet level, but were not essential for the results in Chapter 2, where it was enough for $p(.)$ to be smoothly increasing.

4.2 Theoretical Analysis of the Model

Suppose the gain parameters are fixed for all the routes. Stability, at least on a time-scale of seconds is established by the explicit construction of a Lyapunov function. Let x be the unique stable vector for the system of differential equations (4.1)–(4.2) without time-lags. The rate of convergence can be investigated by linearising this system of differential equations about the stable point. Following the approach in Chapter 2, we obtain

$$\frac{d}{dt}y(t) = -\kappa\Gamma^T\Phi\Gamma(t), \tag{4.3}$$

where

$$\Gamma^T\Phi\Gamma = WX^{-1} + X^{1/2}A^T p'(y)AX^{1/2}. \tag{4.4}$$

In the above expression, the notations have the same definition as before. Without loss of generality, we assume that the capacity of the bottleneck resource is normalised to one. Under steady state conditions, we have for the one-node network

$$y = (|R|w)^{1/N+1},\ p = (|R|w)^{N/N+1},\ x = \frac{(|R|w)^{1/N+1}}{|R|} \text{ and } p' = N(|R|w)^{\frac{N-1}{N+1}}.$$

Suppose further that the network has only two routes. Label the users as r_1 and r_2, respectively. Assume that for user r_i, $i = 1, 2$, the corresponding gain parameters and round-trip times are κ_i and τ_i. Moreover, let $\tau_1 = \tau$ and $\tau_2 = c\tau$, where c is some real positive scalar. Let the users pay a uniform charge of w for using their respective routes. Then $W = wI_2$ and $X = xI_2$, where x is the unique stable point of the system. The network topology is then represented by $A = [1\ \ 1]$. With this set of parameters, the above system reduces to

$$\Gamma^T \Phi \Gamma = \begin{bmatrix} \frac{w}{x} + p'x & p'x \\ p'x & \frac{w}{x} + p'x \end{bmatrix}.$$

In particular,

$$\Gamma^T \Phi \Gamma = (2w)^{N/N+1} \begin{bmatrix} 1 + N/2 & N/2 \\ N/2 & 1 + N/2 \end{bmatrix}, \tag{4.5}$$

with $\phi_1 = (1 + N)(2w)^{N/N+1}$ and $\phi_2 = (2w)^{N/N+1}$. Therefore,

$$\Phi = (2w)^{N/N+1} \begin{bmatrix} 1 + N & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad \Gamma = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.$$

The rate of convergence for this system is determined by the smallest eigenvalue, in this case, $(2w)^{N/N+1}$. The rate of convergence therefore depends on the choice of the willingness-to-pay w, and the threshold for packet marking N. We explore this further in Section 4.5.

We continue with a stochastic perturbation of the linearised equation (4.3) and assume

$$dy(t) = -\kappa \left(\Gamma^T \Phi \Gamma y(t)dt + FdB(t)\right)$$

where F is an arbitrary $|R| \times |I|$ matrix and $B(t) = (B_i(t), i \in I)$ is a collection of independent standard Brownian motions, extended to $-\infty < t < \infty$. We obtain the covariance matrix as $\Sigma = \kappa \Gamma^T [\Gamma F; \Phi] \Gamma$, where $[\Gamma F; \Phi] = \begin{bmatrix} \frac{1}{4(N+1)} & 0 \\ 0 & \frac{1}{4} \end{bmatrix}$, $FF^T = (2w)^{N/N+1} I$ and $\Gamma FF^T \Gamma^T = \begin{bmatrix} \frac{(2w)^{N/N+1}}{2} & 0 \\ 0 & \frac{(2w)^{N/N+1}}{2} \end{bmatrix}$. Let $\kappa = \min(\kappa_1, \kappa_2)$. Then,

$$\Sigma = \frac{\kappa}{8(N+1)} \begin{bmatrix} N+2 & -N \\ -N & N+2 \end{bmatrix}.$$

Thus, the variance for each user is $\sigma_r = \dfrac{\kappa(N+2)}{8(N+1)}$, and the co-variances between users r_1 and r_2 are $\sigma_{12} = \sigma_{21} = -\dfrac{\kappa N}{8(N+1)}$ respectively. The expressions for the variances and co-variances for both users are independent of the charge w, when assumed to be same for both of them. Moreover, a change in the buffer size N for packet marking does affect the magnitude of fluctuations at equilibrium. A higher value of κ results in a higher variance at equilibrium, but allows for a more rapid convergence to equilibrium [45]. In the limiting case, we have

$$\lim_{N\to\infty} \sigma_r = \frac{\kappa}{8}; \quad \lim_{N\to 0} \sigma_r = \frac{\kappa}{4}; \quad \& \quad \lim_{N\to\infty} \sigma_{12} = -\frac{\kappa}{8}; \quad \lim_{N\to 0} \sigma_{12} = 0.$$

Therefore,

$$\frac{\kappa}{8} \le \sigma_r \le \frac{\kappa}{4} \quad \text{and} \quad -\frac{\kappa}{8} \le \sigma_{12} \le 0.$$

Finally, note that the trace of the covariance matrix Σ is $\dfrac{\kappa(N+2)}{4(N+1)}$, and the trace of Φ, the matrix of eigenvalues, is $(N+2)(2w)^{N/(N+1)}$. The above results show that the stability conditions are dependent on κ, irrespective of the threshold for packet marking and the willingness-to-pay w. An interesting observation is that, the lower buffer level at which marking occurs, or the higher statistical variability of traffic at packet level, the *fewer* the possibilities for lag-induced oscillatory behaviour.

Remark 4.1 *The assumption that $\kappa = min(\kappa_1, \kappa_2)$ is appropriate to deduce the approximate bounds for the variability of each user. Suppose the gain parameter is not fixed for each user r. Then $\mathcal{K} = \mathrm{diag}(\kappa_1, \kappa_2)$. Using the linearised system in Section 2.2.2, we have the system (2.18), with*

$$\Gamma^T \Phi \Gamma = (2w)^{N/N+1} \begin{bmatrix} (1+N/2)\kappa_1 & N/2\kappa_1^{1/2}\kappa_2^{1/2} \\ N/2\kappa_1^{1/2}\kappa_2^{1/2} & (1+N/2)\kappa_2 \end{bmatrix}.$$

Since the gain parameters are implicitly involved in the above expression, it is very difficult to get a clean expression for the covariance matrix Σ in terms of the two gain parameters, and to decide anything about its components. However, observe that when $\kappa_1 = \kappa_2 (= \kappa)$, the above expression reduces to the system (4.5) with an extra multiplicative factor of κ.

4.3 Time-Lags

The model with time-lags is clearly a more realistic model to study, specially for multimedia traffic which are prone to delays and loss. In this section, we consider stability on shorter time-scales. Consider the lagged, discrete time system

$$x_r[t+1] = x_r[t] + \kappa_r\Bigg(w_r - x_r[t-\tau_r]\mu[t-\tau_r]\Bigg) \tag{4.6}$$

where

$$\mu[t] = p\Bigg(\sum_{s\in R} x_s[t]\Bigg), \tag{4.7}$$

and τ_r is the round-trip time for user r, $r \in R$. Let the vector x be an equilibrium point of the system (4.6)–(4.7) if $x_r[t] = x_r$, for $t = \dots, 0, 1, 2, \dots$, satisfies these equations.

We follow the approach for the primal algorithm studied in Chapter 2 for a general communication network with $|R|$ users and $|J|$ resources. Assuming the same Lyapunov function $\mathcal{U}(x)$ defined by equation (2.15) in Section 2.2.1 with $|J| = 1$ and the same conditions imposed on $p(.)$, we have the following result:

Theorem 4.1 *The vector maximising the strictly concave function $\mathcal{U}(x)$ is the unique equilibrium point of the system (4.6)–(4.7).*

Proof. The assumptions on $w_r > 0$, and $p(.)$ ensure that $\mathcal{U}(x)$ is strictly concave on $x \geq 0$ with an interior maximum; the maximising value of x is thus unique.

The vector x is an equilibrium point of the system if and only if it solves

$$w_r = x_r p\left(\sum_{s \in R} x_s\right), \quad r \in R. \tag{4.8}$$

Observe that

$$\frac{\partial}{\partial x_r}\mathcal{U}(x) = \frac{w_r}{x_r} - p\Big(\sum_{s \in R} x_s\Big); \tag{4.9}$$

which gives the maximum by setting these derivatives to zero.

Further,

$$\begin{aligned} \frac{d}{dt}\mathcal{U}\big(x(t)\big) &= \sum_{r \in R} \frac{\partial \mathcal{U}}{\partial x_r} \cdot \frac{d}{dt} x_r(t) \\ &= \sum_{r \in R} \left(\frac{\kappa_r}{x_r(t)} \left(w_r - x_r(t) p\Big(\sum_{s \in R} x_s(t)\Big)\right)\right)^2, \end{aligned}$$

establishing that $\mathcal{U}(x(t))$ is strictly increasing in t, unless $x(t) = x$, the unique x maximising $\mathcal{U}(x)$.

The function $\mathcal{U}(x)$ is thus a Lyapunov function for the system (4.6)–(4.7). This is precisely the stationarity condition implied by the partial derivatives of $\mathcal{U}(x)$, a concave function with a unique maximum. This completes the proof.

■

We determine the rate of convergence by considering perturbation about the stable point. The unique equilibrium point of this system does not depend on the round-trip time τ, but its stability does. Suppose $\mu = p(\sum_{s\in R} x_s)$, and suppose p is differentiable at this point, with derivative p'. Let $x_r[t] = x_r + x_r^{1/2} y_r[t]$. We linearise the system (4.6)–(4.7) about x.
From (4.7) we have

$$\mu[t-\tau_r] = p\left(\sum_{s\in R} x_s\right) + \left(\sum_{s\in R} x_s^{1/2} y_s[t-\tau_r]\right) p'\left(\sum_{s\in R} x_s\right). \tag{4.10}$$

For the sake of convenience, we assume $p = p\left(\sum_{s\in R} x_s\right)$, and $p' = p'\left(\sum_{s\in R} x_s\right)$.
Using equation (4.6) we get

$$x_r^{1/2} y_r[t+1] = x_r^{1/2} y_r[t] + \kappa_r w_r - \kappa_r \left(x_r + x_r^{1/2} y_r[t-\tau_r]\right) \mu[t-\tau_r],$$

and using the expression (4.10) for $\mu[t-\tau_r]$ we have

$$\begin{aligned} y_r[t+1] &= y_r[t] - p\kappa_r y_r[t-\tau_r] - \kappa_r x_r^{1/2} p' \sum_{s\in R} x_s^{1/2} y_s[t-\tau_r] \\ &= y_r[t] - \kappa_r \left(\frac{w_r}{x_r} y_r[t-\tau_r] + p' x_r^{1/2} \sum_{s\in R} x_s^{1/2} y_s[t-\tau_r]\right). \end{aligned} \tag{4.11}$$

We can also write the above linearised equation as

$$y_r[t+1] = y_r[t] - \kappa_r \left(\frac{w_r}{x_r} + p' x_r\right) y_r[t-\tau_r] - \kappa_r p' x_r^{1/2} \sum_{\substack{s\in R \\ s\neq r}} x_s^{1/2} y_s[t-\tau_r]. \tag{4.12}$$

Equation (4.12) is the linearised version of equations (4.6)–(4.7) for $|R|$ users and a single bottleneck resource. It is more complicated to represent the system (4.12) in matrix form for some general value of $|R|$ as compared to the models studied in Chapter 2, since the dimension of the matrix will vary according to the largest round-trip time in this system. It is therefore more appropriate to study cases with different round-trip times independently and deduce a general form by induction.

With the present choice of parameters, equation (4.12) becomes

$$\begin{aligned} y_r[t+1] &= y_r[t] - \kappa_r \left((|R|w)^{N/N+1} y_r[t-\tau_r] + \frac{N}{|R|}(|R|w)^{N/N+1} \sum_{s\in R} y_s[t-\tau_r] \right) \\ &= y_r[t] - \frac{\kappa_r}{|R|}(|R|w)^{N/N+1} \left(|R| y_r[t-\tau_r] + N \sum_{s\in R} y_s[t-\tau_r] \right). \end{aligned}$$

Define δ_s, $s \in R$, as

$$\delta_s = \begin{cases} 1 & \text{if } s = r, \\ 0 & \text{otherwise.} \end{cases}$$

Then

$$y_r[t+1] = y_r[t] - \frac{\kappa_r}{|R|}(|R|w)^{N/N+1} \sum_{s\in R} (|R|\delta_s + N) y_s[t-\tau_r]. \tag{4.13}$$

Let $\Upsilon = \max_r \{\tau_r,\ r \in R\}$ be the largest round-trip time of the system. We can rewrite equation (4.13) in the matrix form

$$\begin{pmatrix} y[t+1] \\ y[t] \\ \vdots \\ y[t-\Upsilon+1] \end{pmatrix} = L \begin{pmatrix} y[t] \\ y[t-1] \\ \vdots \\ y[t-\Upsilon] \end{pmatrix} \tag{4.14}$$

where L is a square matrix of order $|R|(\Upsilon+1)$. The matrix L can take the general form

$$L = \begin{pmatrix} I_{|R|} & \cdots & \mathcal{K}L[d] & \cdots & \cdots \\ I_{|R|} & 0 & 0 & \cdots & \cdots \\ 0 & I_{|R|} & \cdots & \cdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & I_{|R|} & 0 \end{pmatrix}$$

where $L[d]$ is a square matrix of order $|R|$ a null matrix of order $|R|$, $I_{|R|}$ is the identity matrix of order $|R|$, and $\mathcal{K} = \text{diag}(\kappa_r,\ r \in R)$ is a diagonal matrix of the gain parameters.

It is interesting to note that the structure of L is not obtained as in equation (2.39), because the algorithm considered for the single scarce resource is a variant of the primal time-lag model considered in Section 2.2.4, equation (2.35). In fact, in the present model, we consider a round-trip time of τ_r, which is the total time taken from the end-user r to the resource and back.

In the matrix L, only the first row contains information about the system's behaviour. The equilibrium point x of the system (4.6)–(4.7) is stable if L has spectral radius less than unity. The spectral radius depends not only on the predetermined contract parameters R, w, and N, of the model, but also on the gain parameter κ_r for each user r, $r \in R$. While the predetermined parameters are known and fixed, assuming the resource has finite capacity, we do not have much information about the choice of κ_r.

In the following sections, we study the particular scenario when two routes pass through a single bottleneck resource. We determine bounds for the gain parameters to maintain stability of the system. We also study how the variation of the round-trip time influences stability and the bounds of the corresponding gain parameters. The result will provide the end-users with a range of gain parameters to choose from, in order to make efficient use of the network and to maintain stability at equilibrium.

4.4 Stability Conditions

Before we proceed to analyse the single node system numerically, it is appropriate to mention some bounds for stability. The stability conditions were derived using the concept of retarded differential equations (see [23, 51, 67, 93] for more details). Suppose x_r is an equilibrium point of the system (4.1)–(4.2). Let $x_r(t) = x_r = u(t)$, and write p, p' for the values of the functions $p(.)$, $p'(.)$ at x_r. Observe that x_r and p are related by w/p, where we assume $w < C$, the capacity of the resource. Then, linearising the system (4.1)–(4.2) about x_r, we obtain

$$\frac{d}{dt}u(t) = -\alpha u(t-\tau),$$

with $\alpha = \kappa_r(p + x_r p')$. The equilibrium point of the differential equation (4.1) is stable if,

Condition I:

$$\kappa_r(p + p'x) < \frac{\pi}{2\tau_r}.$$

If we consider the discrete form of the differential equation, then the system is locally stable if,

Condition II:

$$\kappa_r(p + p'x) < 2\sin\left(\frac{\pi}{2(2\tau_r + 1)}\right).$$

Apart from the stability in Condition I, the local convergence is non-oscillatory if,

Condition III:

$$\kappa_r(p + p'x) < \frac{1}{e\tau_r}.$$

Table 4.1: Upper bounds for gain parameters.

τ	**Condition I**	**Condition II**	**Condition III**	**Spectral radius = 1**
1	0.3498	0.2228	0.0820	0.2223
5	0.0700	0.0634	0.0164	0.0633
10	0.0350	0.0333	0.0082	0.0333
50	0.0070	0.0070	0.0016	0.0070

Since the model we analyse is a discrete time-lag one, the stability condition II should be sufficient.

Suppose the round-trip times for both users are the same and equal to τ. Table 4.1 gives the upper bounds for the gain parameters of both users obtained from the above conditions for different round-trip times. The last column gives the critical gain value for the time-lag system when the spectral radius, denoted by ϱ, of the linearised system is unity. This is discussed in more detail in the following section. Two observations can be made from the table: firstly, the numerical values in the last column compare best with those of condition II; and secondly, as the round-trip time increases, all the values converge to the same quantity, with the exception of the values corresponding to condition III.

If we assume the first condition for local stability discussed above, then for the continuous case, we should have,

$$\frac{\kappa_r p(N+1)}{2} < \frac{\pi}{2\tau_r}$$

and if τ_r is fixed as τ say, and noting that $\tau_2 = c\tau_1$, then we have the upper bounds for the gain parameters for user r_1 and user r_2 as

$$\kappa_1 < \frac{\pi}{p\tau(N+1)}, \quad \kappa_2 < \frac{\pi}{pc\tau(N+1)}.$$

Let $l = \dfrac{\pi}{p(N+1)}$, which is fixed for a defined set of parameters. Then

$$\kappa_1 < \frac{l}{\tau}; \quad \kappa_2 < \frac{l}{c\tau}.$$

Thus, the relationship between κ_1 and κ_2 depends on c. Note that c is chosen such that the round-trip times take integral values. If $c > 1 \implies \kappa_2 < \kappa_1$; if $c < 1 \implies \kappa_1 < \kappa_2$; if $c = 1$, then both are bounded by $\dfrac{l}{\tau}$. Therefore $\dfrac{l}{\tau} \left(= \dfrac{\pi}{\tau(N+1)(|R|w)^{N/(N+1)}} \right)$ can be taken as an upper bound for both gain parameters. An increase in the number of users using the single resource will cause more competition for bandwidth, and hence, a smaller range for the gain parameters. The system is therefore locally stable under the above conditions.

Since we are analysing the discrete time-lag algorithm, then condition II is more appropriate. In this case $l = \dfrac{4}{p(N+1)} \sin\left(\dfrac{\pi}{2(2\tau+1)}\right)$. As the round-trip time increases, the same bounds, as with condition I, are obtained for the gain parameters.

We present some numerical results for the discrete time-lag model in the next section.

4.5 Numerical Results

For a single bottleneck resource with two users, equation (4.13) reduces to

$$y_r[t+1] = y_r[t] - \frac{\kappa_r}{2}(2w)^{N/N+1} \sum_{s=1,\,2} (2\delta_s + N) y_s[t - \tau_r], \; r = 1,\, 2. \tag{4.15}$$

We study the stability of the system under different scenarios. The spectral radius is computed for each scenario, and its variation with the gain parameters studied in order to maintain stability. In this section, we investigate mainly the variation of the round-trip times with the gain parameters.

The experiment was carried out in two parts. In the first stage, the round-trip times for both users were kept equal and the bounds for the gain parameters over which the system remains stable obtained. In the second part, we study the situation when the round-trip times are different. In the sequel, we denote the spectral radius of the matrix L by ϱ.

In this model, we take $|R| = 2$ and the single-node capacity is assumed to be unity. It is also assumed that all the users pay a uniform charge of w for using their respective routes. Since $p(.)$ is the probability for packet marking, we must have $p(y) \leq 1$, so that $w \leq \dfrac{1}{|R|}$ or $w \leq \dfrac{1}{2}$ for the two-user case. This is the condition for the dynamical system to be feasible.

For the numerical study of this system, we choose $N = 10$ and $w = 0.4$. This gives us the following set of values at equilibrium for our study: $y = 0.97992$, $x = 0.48996$, $p = 0.81639$, and $p' = 8.33125$.

4.5.1 Round-trip time same for both sources

In the first instance, we keep the round-trip times the same for both users, and study the stability as the gain parameters vary. For the matrix L to be convergent, we must have its spectral radius less than unity. The smaller the spectral radius, the more stable is the system. The range over which the gain parameters for both users vary will describe a region, each point within which keeps the spectral radius less than unity. We shall call the feasible region the *stability region* for the system. Any pair of gain parameters (κ_1, κ_2) outside this region will lead to an explosive state. Hence, for routes with known round-trip times, each source can choose, among a set of values, the appropriate quantity to maintain stability while having an optimal and fair share of the bandwidth throughout its session. Moreover, the choice of the parameters for this particular model makes the gain parameters symmetric when the round-trip times are same for the two users. This implies that studying the behaviour of the gain parameter for only one user is sufficient to understand the system. In this section, we shall assume $\kappa\,(= \kappa_1 = \kappa_2)$ to denote the gains for both sources.

When both round-trip times are the same, say τ, following the results from Section 4.3, we can denote the L matrix of order $2(\tau+1)$ in block form by L_τ where

$$L_\tau = \begin{pmatrix} I & 0 & \cdots & 0 & \mathcal{K}L[\tau] \\ I & 0 & \cdots & \cdots & 0 \\ 0 & \ddots & \cdots & \cdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & I & 0 \end{pmatrix}. \tag{4.16}$$

In the above expression, we assume the following: $\mathcal{K} = \mathrm{diag}(\kappa_1,\, \kappa_2)$, I is the identity matrix of order two, and

$$L[\tau] = \begin{pmatrix} \beta & \alpha \\ \alpha & \beta \end{pmatrix},$$

where $\alpha = -\dfrac{N}{2}(2w)^{N/N+1}$ and $\beta = -\dfrac{N+2}{2}(2w)^{N/N+1}$. The remaining blocks in L_τ are null matrices of order two. We present the results for $\tau = 1$, 5 and 10 only. The conclusions can be generalised and will hold for integral values of τ.

Computational results

The general variation of the spectral radius with the gain parameters is first studied to identify the feasible region for the system. The two graphs in Figure 4.1 show this relationship for $\tau = 1$. The stability region is the space in Figure 4.1a, enclosed by the axes and below the mesh when ϱ is less than unity. In the contour plot, the darker region shows that portion of the plot when $\varrho < 1$. Note that, in Figure 4.1b the horizontal axis corresponds to κ_1, while the vertical axis corresponds to κ_2. This suggests that the system is stable for any pair of (κ_1, κ_2) lying in the darker region. In this case, κ can be safely taken to be less than 0.2. A similar behaviour is observed in the cases when $\tau = 5$ and 10. For a round-trip time of five, the appropriate range of κ to maintain stability is $(0, 0.06]$ while it is $(0, 0.03]$ when $\tau = 10$. The experiment suggests an inverse relationship between the gain parameters and the round-trip times in the system.

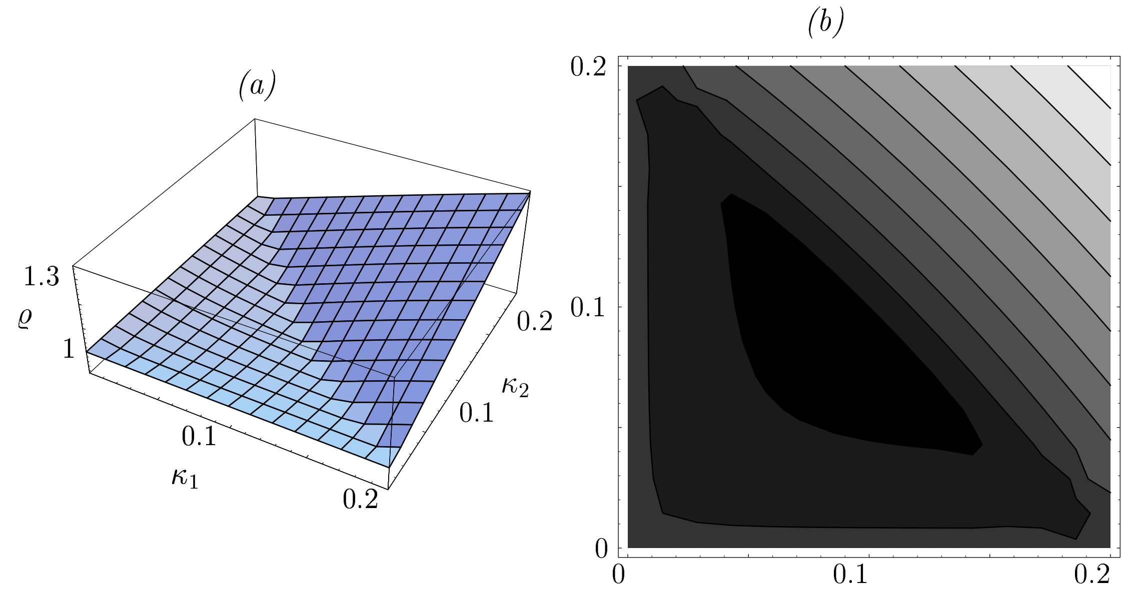

Figure 4.1: Graph (a) shows the three-dimensional plot of the spectral radius with the gain parameters for $\tau = 1$. The stability region is enclosed by the axes and below the plane $\varrho = 1$. Plot (b) is the contour graph of the same scenario. The darker portions represent the stability regions. The horizontal axis corresponds to κ_1 and the vertical axis to κ_2 respectively.

For a fixed value of κ_1 (or κ_2), Figure 4.2a shows the relationship between the spectral radius of the matrix L and κ_2 (or κ_1). The left panel shows the general variation of ϱ as κ_2 (or κ_1) increases. We are mainly interested in the portion of the curve below $\varrho = 1$ to study stability. This region is represented in Figure 4.2b. It can be observed that as κ increases (from within a neighbourhood of zero), the spectral radius decreases from unity. Let κ_l denote the lower bound for the gain parameters corresponding to which $\varrho < 1$. The spectral radius continues its descent to reach a minimum when $\kappa = \kappa_{min}$. Then, it increases monotonically to cross the *stability limit* at $\kappa = \kappa^*$, the upper bound for the gain parameters (the critical gain). The open interval (κ_l, κ^*) gives the range of κ for the system to remain stable. In the present context, we shall call this interval the *stability interval.* For the different round-trip times studied, the stability interval for the gain parameters have lower bounds zero and upper bounds as per the last column in Table 4.1. In general, we deduce that for the system to be stable, κ should neither be too small nor too large (that is, not too close to the boundary) to ensure stability, as indicated by the darkest region of Figure 4.1b. The nearer the gain to the boundary, the larger is the number of oscillations in the system, with increasing amplitudes. This situation is not desirable. In Section 4.6, we explore this phenomenon using convergence times. Furthermore, as τ increases, the stability interval (κ_l, κ^*) shortens until the system is no longer stable for extremely large round-trip times.

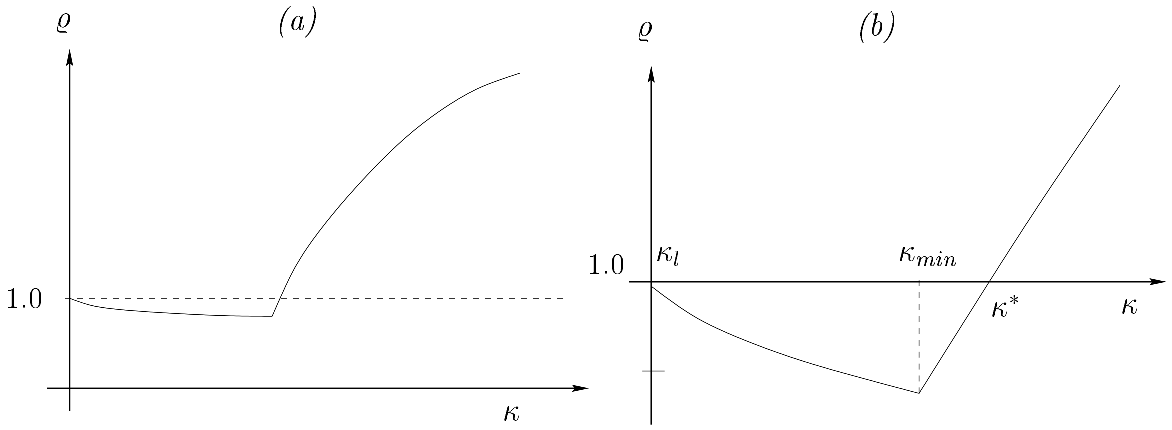

Figure 4.2: General relationship between spectral radius and gain parameter.

The variation of the spectral radius with κ_1 (or κ_2) for different values of κ_2 (or κ_1) is shown in Figure 4.3a. The topmost curve is when $\kappa_2 = 0.2$ and the curve nearest to the line $\varrho = 1$ is for $\kappa_2 = 0.02$. It is interesting to note that on a microscopic scale, when ϱ increases from the minimum point κ_{min}, the curves have steep ascents for all values of κ_2 (or κ_1). Moreover, as κ_2 (or κ_1) increases, the range of κ_1 (or κ_2) for which $\varrho < 1$ decreases up to $\kappa_2 \approx 0.2$ when the system is no longer stable. We can also observe that the lower envelope of the family of curves for each κ_2 is convex. A similar behaviour is noticed when $\tau = 5$ and 10 over smaller stability intervals.

Next, we study the variation between the gain parameters at the stability limit. Figure 4.4a shows this situation for $\tau = 1$, 5 and 10. The open space bounded by the curves and the axes represent the stability region for the single link network. The

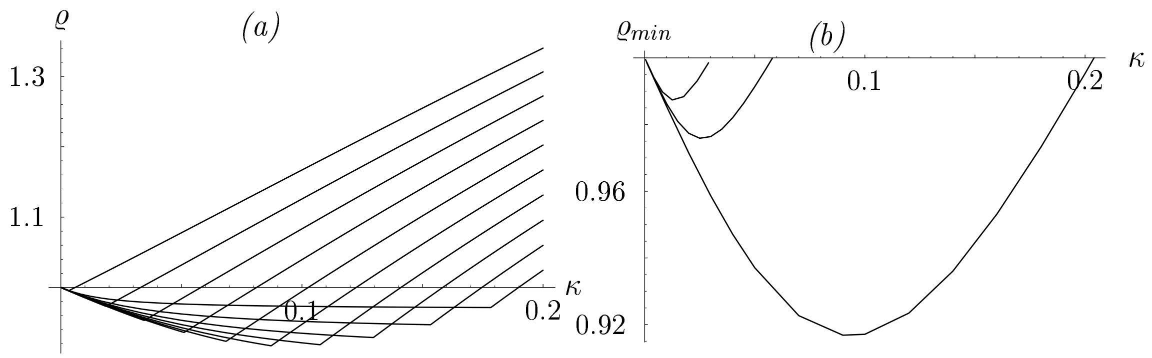

Figure 4.3: Plot (a) shows the variation of the spectral radius with κ_2, for different values of κ_1, and vice-versa, for $\tau = 1$. The second plot shows the variation of the minimum spectral radius with κ for $\tau = 1, 5, 10$. The smallest curve corresponds to $\tau = 10$, and the largest curve to $\tau = 1$.

curves are defined by κ^*, and any pair of gain parameters will maintain stability if they satisfy the inequality $\kappa_1 + \kappa_2 \leq \kappa^*$. The larger the round-trip time, the smaller the ranges of κ over which the system is stable, and hence the smaller the stability region. The boundary of the stability curve is a concave function with no feasible points on the boundary. The concavity becomes more obvious as the round-trip time increases. This result is supported by simulation.

The variation between the minimal spectral radius and the corresponding gain parameter, κ_{min}, in each curve of Figure 4.3a, is shown in Figure 4.3b. Firstly, we confirm our hypothesis in the first plot that the minimum spectral radius for each fixed κ_2 (or κ_1) within the stability region is a convex function (approaching a parabolic path), obtained as the envelope of the family of curves in Figure 4.3a. The second plot in Figure 4.4 shows the relationship between κ_{min} for both users. Again, as in the plot for the stability region, the upper curve corresponds to the case when the round-trip time is unity, and the lowest curve corresponds to a round-trip time of 10. Further, observe that the curves in Figure 4.4 are inverses of each other. For any communication system, a smaller spectral radius is desirable in an optimal system at equilibrium. For each round-trip time, the loop created by the corresponding curves in Figure 4.4

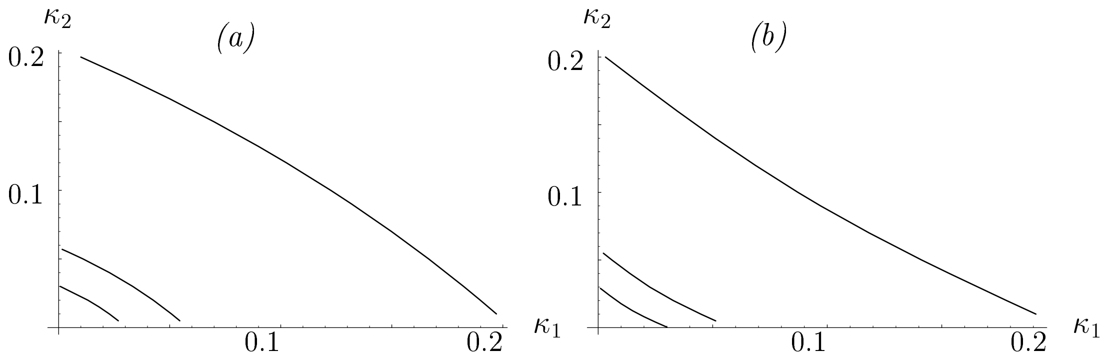

Figure 4.4: Plot (a) shows the stability regions for $\tau = 1$, 5, 10 defined by the stability limits and the gain parameters. The uppermost curve corresponds to $\tau = 1$, while the lowest one corresponds to $\tau = 10$. The second plot refers to more restricted stability regions identified by the gain parameters at the minimum spectral radii. In these regions the oscillations are less than those in plot (a).

defines the region of highest oscillations in the system. A more stable feasible region, described by the curves in Figure 4.4b, is therefore preferred since there will be less oscillations when the gains are chosen from this area.

4.5.2 Different round-trip times for both users

In this section, we study the more general case when the round-trip times for both users are different. As in the previous section, we shall identify the stability regions for various scenarios by using the convergence property of the matrix. This case is more complex due to the lack of symmetry in the system. Thus, the stability intervals for both users are not equal.

Let the round-trip time for user r_1 be τ_1 and that for user r_2 be τ_2. Let $\tau = \max\{\tau_1, \tau_2\}$. Since the dimension of the matrix L depends on the maximum round-trip time, we have the dimension of the array as $2(\tau + 1)$. If we partition the matrix L in blocks of order two, we can represent the general structure of the matrix L for the two-user case when the round-trip times are different as

$$L_\tau = \begin{pmatrix} I & \cdots & L[\min\{\tau_1,\tau_2\}] & 0 & \cdots & \mathcal{K}L[\tau] \\ I & 0 & \cdots & \cdots & \cdots & 0 \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ 0 & \cdots & \cdots & 0 & I & 0 \end{pmatrix}. \tag{4.17}$$

In this case

$$L[\tau] = \begin{pmatrix} 0 & 0 \\ \alpha & \beta \end{pmatrix}, \quad L[\min\{\tau_1,\tau_2\}] = \begin{pmatrix} \beta & \alpha \\ 0 & 0 \end{pmatrix},$$

and α, β have the same definition as in Section 4.5.1. We present the results only for the case when $\tau_1 = 1$ and $\tau_2 = 5$. Other combinations give similar results on different scales.

Computational results when τ_1 and τ_2 are different

Consider the case when $\tau_1 = 1$, $\tau_2 = 5$. The general variation of the spectral radius with the gain parameters is studied to identify the feasible region, that is, the region bounded by the coordinate axes and the plane $\varrho = 1$. The behaviour of the system is shown in Figure 4.5. It can be observed that the first plot is not symmetric as compared to Figure 4.1a. In fact, the range for κ_1 is longer than that for κ_2, in relation to their round-trip times, as predicted in the previous section. This is shown by the third contour from the centre of the contour plot in the same figure. Further the stability region has a maximum point moving parallel to the κ_2 axis. For this model, the stability interval is not fixed. Interestingly, the second user has a wider interval as compared to the case when the round-trip times were the same. That for the first user, however, is significantly reduced. Therefore, when the round-trip times are different, the algorithm considered favours the route with the longer round-trip time while maintaining optimality. Similar observations are made for the other

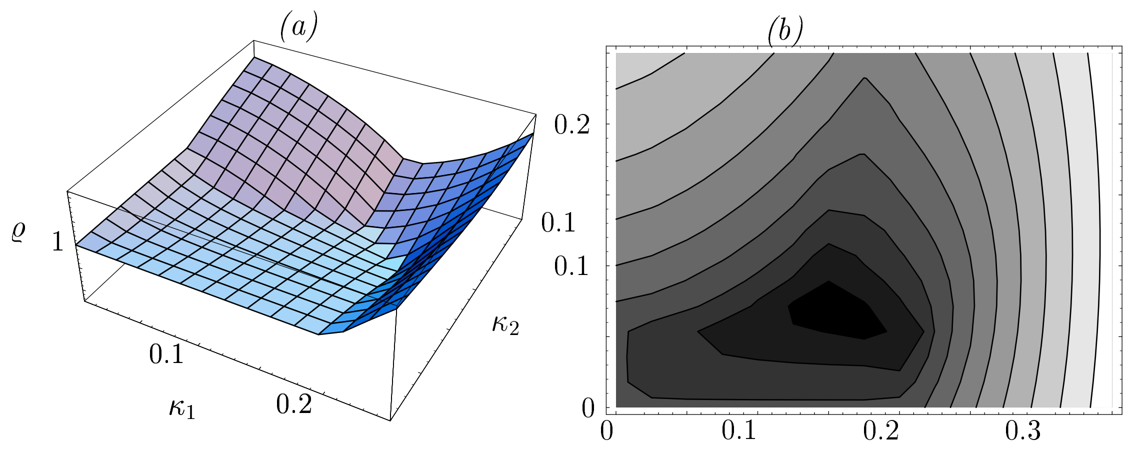

Figure 4.5: Plot (a) shows the three-dimensional plot of the spectral radius with the gain parameters for $\tau_1 = 1$, $\tau_2 = 5$. The stability region is enclosed by the axes and below the plane $\varrho = 1$. Plot (b) is the contour graph for the same scenario. The darker portions provide more stability. Note that the horizontal axis corresponds to κ_1, and the vertical axis to κ_2 respectively.

scenarios studied under different scale factors.

Next, if we consider slices of the three dimensional graph in Figure 4.5a, we have different shapes for different values of κ_1 or κ_2. The family of curves showing the variation of the spectral radius with κ_1, for different values of κ_2 is presented in Figure 4.6a. The uppermost curve corresponds to $\kappa_2 = 0.15$, and the lowest curve is for $\kappa_2 = 0.07$. The last curve having a feasible point corresponds to $\kappa_2 = 0.15$. It can be observed that the lower bounds of some intervals are not necessarily zero as obtained in the symmetric case.

A spectral radius of unity is an upper limit for the stability of the system. The left hand plot of Figure 4.7 shows the variation of the gain parameters for the two users when the spectral radius is unity. We observed earlier that for this model there are two non-zero values of κ_1 for each fixed value of κ_2 where the spectral radius is unity, κ_l and κ^* respectively. The region enclosed by the curve and the axes is the stability region for the single node system. Its geometry conforms with the contours in Figure 4.5b. Further, we note that the intercepts are the upper bounds for the

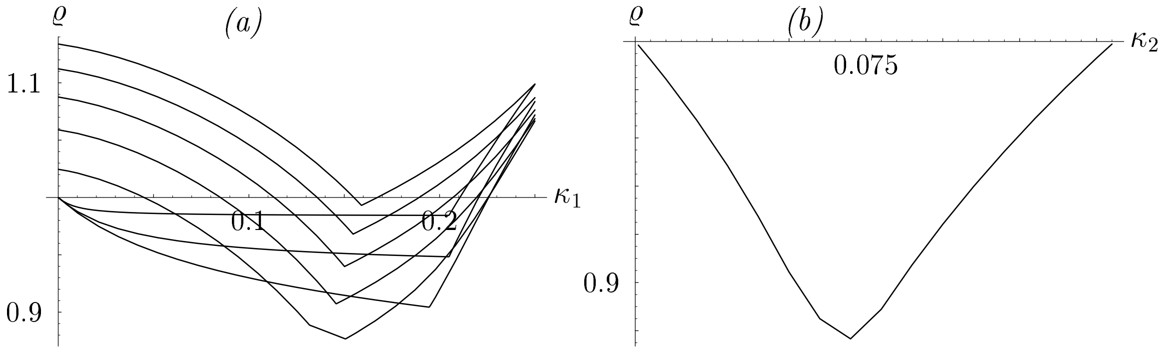

Figure 4.6: Plot (a) shows the variation of the spectral radius with one gain parameter, for different values of the other gain parameter, and for $\tau_1 = 1$, $\tau_2 = 5$. The second plot represents the envelope of the family of curves in plot (a).

stability intervals for the two gain parameters. A more peculiar behaviour is the maximum point (in this case (0.161,0.155)) of this curve with respect to the κ_1-axis. This general behaviour is the same for other sets of round-trip times studied and is verified by simulation.

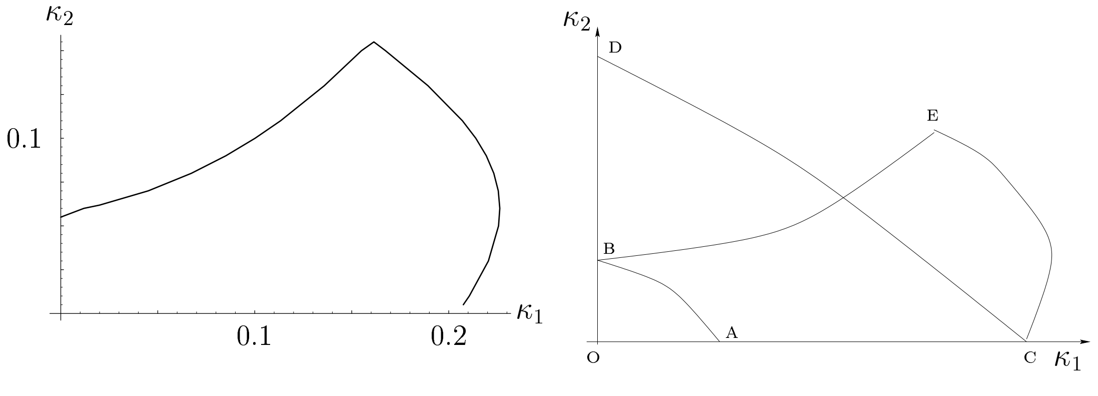

Figure 4.7: Stability regions when $\tau_1 = 1$ and $\tau_2 = 5$.

A closer look at Figure 4.6a reveals a particular pattern for the minimum of each curve in this family. This behaviour is explained from the plot in Figure 4.6b which gives the variation of the minimum spectral radius with κ_2. For this model, the curve is symmetrical about the line $\kappa_2 = 0.07$ when the curve attains its minimum value of $\varrho = 0.877$. This value of the spectral radius is smaller than the values of 0.976 and 0.917 obtained separately for $\tau = 5$ and $\tau = 1$, as shown in Figure 4.3b. A smaller value of ϱ gives a highly stable system.

In the right hand panel of Figure 4.7, the curves AB and CD identify the stability regions when both round-trip times take equal values of 1 and 5 respectively. Curve CEB in the same figure, is simply the left hand graph in Figure 4.7. The superimposed curves clearly show the increased stability interval for the second user, corresponding to a larger round-trip time. In fact, the stability region $OCEB$ is about ten times bigger than that of OAB, but about 3% less than that of OCD.

Thus, having routes with different round-trip times are likely to provide a more stable system. This joint effect of having different round-trip times is more advantageous for the user with longer round-trip time as its stability interval for the gain parameters increases, allowing more choice.

Remark 4.2 *The choice of parameters in this chapter represents just one of the multiple scenarios which can occur in reality. Varying the predetermined parameters has consequent effects on the stability regions and hence on the flexibility for the end-users to choose the gain parameters. The analysis of this work depends mainly on the willingness-to-pay* w*, by the user, the threshold* N*, for packet marking and the number of users for the scarce resource. The choice of the function* $p(.)$*, a function of* $R, N,$ *and* w*, also influences the choice of the gain parameters. A user may pay a charge to receive a flow proportional to this willingness-to-pay. An increase in this quantity decreases the size of the feasible region and hence the limits of the stability interval. Experimentally, we have observed that the willingness to pay for a user is inversely proportional to the upper bound of the gain parameter. Further, the threshold for packet marking can also be changed by the routers depending on the traffic. Increasing* N *has the general effect of compressing the stability region. Thus, in a dynamic environment, these parameters can be properly adjusted to allow a win-win policy for* *both the resource and the users.*

4.6 Convergence Times

In the previous section, we have seen how the stability of the system depends on the round-trip time and the gain parameter of each connection. We shall now explore the effect of the gain parameters and the round-trip times on the convergence times for the algorithm (4.1). In Chapter 2, we have proved that the rate of convergence to equilibrium increases with the gain parameter at the cost of more variability. Here, we also deduce a possible relationship between the convergence times and the spectral radius of the matrix L derived earlier.

The convergence times are obtained by simulating the deterministic discrete time-lag algorithm for a long period of time until the equilibrium point has been reached, and the system oscillates within an $\epsilon-$neighbourhood of the stable point. For the sake of convenience, we assume discrete time increments which can be scaled to suit any real-life model. For example, we may choose one time unit to represent 200msec for the TCP/IP. Further, all the initial flows are assumed to start from zero. Let the convergence times for users r_1 and r_2 be T_1 and T_2 respectively. We perform the experiment when the round-trip times are equal as well as when they are different. It can be noted that, for equal round-trip times, it is sufficient to study the behaviour of one user to understand the system's behaviour.

Let the gain parameter $\kappa_2 = \kappa'$. The convergence times for the rates of flow for both users are determined as the gain parameter for the first user r_1 is varied. Figure 4.8 shows the general behaviour of the convergence time with the gain parameter. In general, when $\kappa_1 = \kappa_l$ is in the neighbourhood of zero, the convergence times for both users are very long but without oscillations. As κ_1 increases, the convergence times for both users decrease to a certain minimum point. The times then grow rapidly with increasing and large oscillations until explosion. This aspect has been highlighted in the analysis of the primal and dual algorithms for the random network. The behaviour

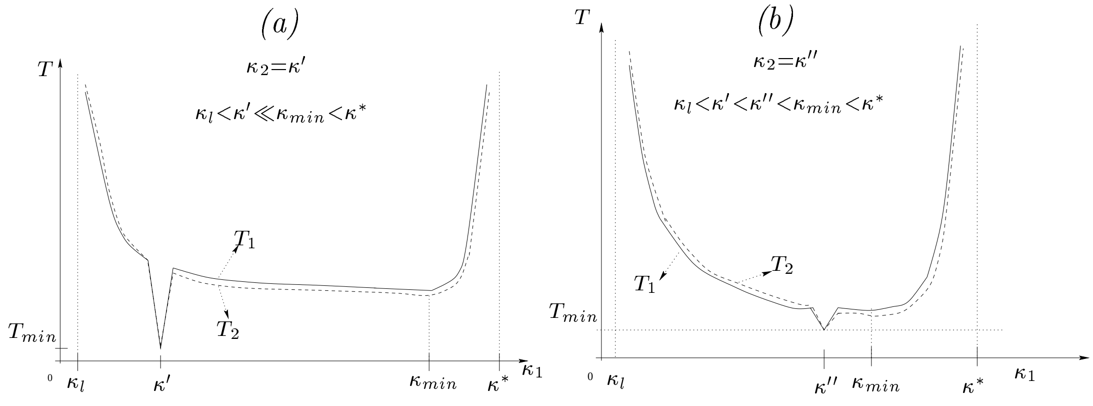

Figure 4.8: Variation of convergence times with κ_1 for $\kappa_2 = \kappa', \kappa''$ and $\kappa' < \kappa''$.

of the convergence time is convex and the flatness of the curve is determined by the value of κ_2. The convexity increases with κ_2.

Next, the variation of the convergence times for the flow rates for the second user with κ_1 is observed. When $\kappa_1 \neq \kappa_2$, the route with larger κ overshoots the stable point and converges from above, in contrast to the route with smaller κ which converges from below. Further, the route with larger gain oscillates more before settling down to the equilibrium point. Suppose $\kappa_1 > \kappa_2$, then x_1 will have larger amplitude and converge from above to the stable point with a high degree of jittering. As κ_1 decreases (but still larger than κ_2), the number of oscillations also decreases. For $\kappa_1 = \kappa_2$, the curves for the flow rates of both users behave identically: they oscillate together and converge together to the same equilibrium point. At any point of time both rates are of the same value. When κ_1 is decreased further so that $\kappa_1 < \kappa_2$, the roles are interchanged: the curve corresponding to x_2 now converges from above and that of x_1 corresponding to κ_1 converge from below. The larger the difference in gains, the more oscillations the upper curve shows. These aspects can be depicted from Figure 4.8a and 4.8b corresponding to $\kappa_2 = \kappa'$ and $\kappa_2 = \kappa''$ with $\kappa' < \kappa''$.

An interesting feature is that the range over which the system converges agrees with the stability interval for the gain parameters. If we compare the behaviour of the convergence time with the spectral radius analysis in the previous section, we highlight some common features. The minimum convergence time is achieved at κ_{min}. From this point, the convergence time increases sharply as does the spectral radius (compare Figure 4.8 with Figure 4.2a) until the stability limit is reached at $\kappa = \kappa^*$ and explosion occurs. When $\kappa_1 = \kappa_l$, the spectral radius is in the neighbourhood of unity, and in this region, the convergence time is again large but without oscillations. Thus, the convergence time is proportional to the spectral radius: the smaller the spectral radius, the faster the convergence with a lesser number of oscillations, and the more stable is the system.

Another important observation in this particular model is that when $\kappa_1 = \kappa_2 = \kappa'$ (Figure 4.8a), there is a sudden drop in the convergence times (this facet is not depicted in the spectral radius analysis). The drop is very sharp when $\kappa' \ll \kappa_{min}$. As κ' approaches κ_{min}, the drop increases, to become indiscernible when $\kappa' \geq \kappa_{min}$ (see Figure 4.8b where $\kappa'' > \kappa'$). The feature has the advantage of allowing the users to choose their gains to be equal and less than κ_{min}, provided it is known, thereby getting a very fast convergence. Otherwise, the convergence times will depend only on the minimum spectral radius for fastest convergence.

We can compare the behaviour of the convergence times with that of the stability region. By determining the convergence times for points within the feasible region, we observe that the minimum convergence times fall within the darkest region of the contour graph in Figure 4.1. Also in this region, the spectral radii are minimum. This strongly supports the proportionality criterion between the two variables. Further, the convergence times are largest near the boundaries. Another interesting point is that along the curve in Figure 4.4b, the convergence time is small, with either little or no oscillations. Any point within the loop above this region causes large oscillations to the system. Hence, if we want a highly stable system, the users should first try to select their gains equal, within the new improved stability region. If this is not possible, then any point within the reduced stability region will maintain stability.

Figure 4.9 summarises the above discussion when the round-trip times are equal.

Assuming the users are cooperative, it is best for their chosen gains to fall along the line segment LM (where $\kappa_1 = \kappa_2$). This will provide very fast convergence to

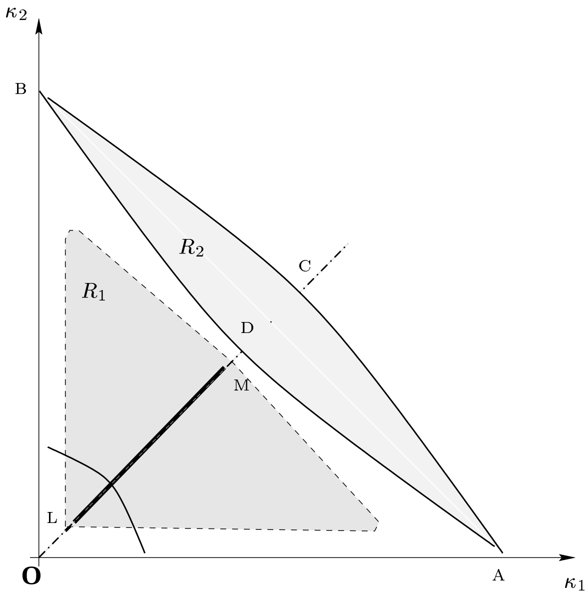

Figure 4.9: Stability region for the gain parameters when the round-trip times are equal.

equilibrium, and a highly stable system. If this is not possible, then the second priority to maintain stability with very little or no fluctuations, it is appropriate that the users choose their gains to fall within the region R_1. In case their choice falls within the loop[1] R_2, then there is bound to be an increasing number of oscillations with large amplitudes, and the convergence time will be very long. The same situation is observed in case the users' choice is outside the shaded region R_1, but within the region OAB. Any point outside this region is unacceptable if we want the system to be stable. Thus, assuming the round-trip times are known, a user can be offered the choice of his corresponding gain parameter interactively with another user sharing a route with identical round-trip time to maintain a high stability.

Repeating this experiment for a series of equal round-trip times has revealed the same

[1]Note that the two curves ACB and ABD are the topmost curves of Figure 4.4a and 4.4b, and are inverses of each other.

conclusions. However, the larger the round-trip times, the longer the convergence times.

When the round-trip times are not equal, we have a similar behaviour as in the previous case with the exception, that there are no visible drops under whatsoever circumstances. This may be due to the asymmetry of the system. Again, the minimum convergence times are obtained in the darkest portion of the contour plot in Figure 4.5, and vary with the spectral radius. The users are encouraged to choose their gains within this region for maximum stability of the system.

In this chapter, we have investigated a network with a single bottleneck resource. The work presented is a further step in the exploration of how the end-users can choose the model parameters so as to make efficient use of the network. In particular, we have shown how the end-nodes may choose the gain parameters, assuming the round-trip times of the different routes are known, in order to maintain the stability of the system while having a fair share of the bandwidth. In the experiment, we have tried to combine the benefits of a large gain parameter, and a small spectral radius, to control the stability of the system. Further work in this area is to generalise this model for R users and determine how the users may choose the gains to minimise the spectral radius of the L matrix.

Chapter 5

Simulations

In the previous chapters we have investigated a class of rate control algorithm, and obtained a series of conditions for the stability of the system. The study will be incomplete without the simulation of the algorithms in different situations. In this chapter, we turn our attention to the simulation aspect. In Chapter 3, we have discussed the choice of a network to test our results. This is important but also difficult due to the unpredictable speed at which the communication networks are expanding.

In the sections that follow, the deterministic and the stochastic cases of the primal and the dual algorithms, discussed in Chapter 2, are simulated for a random network. We also present some simulation results for a TCP-like algorithm and its generalisation. In our approach, we consider two versions of the primal algorithm, namely congestion indication with joint feedback and with individual feedback. On the other hand, the dual algorithm is also simulated for the regular network described in Chapter 3. We show that the results can also hold for networks with definite structures.

For the major part of the simulation, we choose a network which has a random structure described by the matrix A, whose elements are taken as independent Bernoulli random variables, each taking the value 1 with probability ρ $(0 < \rho < 1)$ and 0 otherwise. A larger ρ implies a more connected network. This is the same network topology used in the analysis of the rate control algorithms in Chapter 3. The random network can also be generated from other distributions.

The elements of the delay matrix d for this random network are also taken as independent random variables uniformly distributed over the set $\{0, 1, \ldots, D\}$. Let $w = (w_r = \sum_{j \in J} A(j, r), r \in R)$ and let $V = (v_j \sum_{s: j \in s} A(j, s), j \in J)$. The vector w gives the length of the routes, that is, the number of resources through which each route passes, for each source in the network. Similarly, the vector V gives the number of routes which passes through each node. The choice of parameters is convenient since, under these assumptions, we get a unique stable solution $x_r = 1$, $r \in R$ and $\mu_j = 1$, $j \in J$ as desired for the systems (2.20)–(2.21) and (2.40)–(2.41) considered in Chapter 2.

In this chapter, the following parameter choices were assumed unless otherwise specified. We choose a network having one hundred resources ($|J| = 100$) and one thousand routes ($|R| = 1000$). The maximum time delay D between a source and a resource in the network is chosen as 10 and $\rho = 0.1$. The choice of $D = 10$ causes the round-trip delay to be 20 as will be observed in the graphs later on. This observation may suggest that the gain parameter κ cannot exceed $1/D$ as studied in Chapters 3 and 4 respectively. The particular choice of ρ means that on average 100 routes will pass through one link of the network and 10 resources will lie on one route.

All the simulations were performed for the discrete time-lag systems, in a C environment. The gain parameters are assumed fixed for all the users throughout the simulation. The fixed gain parameter κ and the simulation time t are varied and the data (rate or shadow prices) for the routes and links, unless mentioned otherwise, plotted with time in order to be able to closely follow the various paths. It has been verified that the choice of the links/routes in obtaining the various graphs is immaterial. This is because any randomly chosen subset of sources R and resources J of the network show the same characteristics. For simplicity, we have chosen our time scales to be discrete. The time horizon can be scaled to meet any dimension where the algorithms are being applied.

The deterministic systems for each model considered is simulated over a long time period for different values of κ. The objective is to obtain a range for κ for which the system is stable. The system becomes unstable when the gain parameter is greater than its critical gain, κ^*. Recall that the critical gain parameter provides an upper bound for convergence when the spectral radius of the time-lag system becomes unity. In fact, for the system to converge to the equilibrium value of unity for both x and μ and be stable, we must obtain the range for κ such that it neither takes too much time to converge nor does it diverge. Neither of these is desirable. Instead, we prefer a trade-off between the speed of convergence and the magnitude of oscillations at equilibrium. In the present context, we say that the flow rates or the shadow prices have converged if they oscillate within an ϵ-neighbourhood of the equilibrium point.

Using the gain parameters obtained from the deterministic algorithm, we then simulate the stochastic versions and observe the behaviour of the flow rates and the shadow prices at equilibrium.

5.1 The Primal Rate Control Algorithm

Consider the system (2.20)–(2.21) in Section 2.2. Two versions of this model have been considered for simulation: congestion with joint feedback, and congestion with individual feedback. The deterministic models give the stability interval for the gain parameters. In the stochastic models, the feedbacks are generated as a continuous stream of Poisson[1] random variables.

5.1.1 Congestion indication with joint feedback

We first consider the following discrete algorithm

$$x_r[t+1] = x_r[t] + \kappa\Big(w_r - x_r[t]\sum_{j\in J} A_{jr}\mu_j[t-d(j,r)]\Big) \tag{5.1}$$

[1]the generation of a Poisson random variable is explained in [29].

where $\mu_j[t] = \mu_j[0]$ if $t < 0$ and

$$\mu_j[t] = \sum_{s \in R} A_{jr} x_s[t - d(j,s)]/v_j, \tag{5.2}$$

with $x_s[t] = x_s[0]$ if $t < 0$. This is in fact the time-lag system (2.35)–(2.36) discussed in Chapter 2 with $p_j(y) = y/v_j$. In this model, the flow rate $x_r[t]$ increases linearly by an amount κw_r and decreases multiplicatively at a rate proportional to $\kappa x_r[t]$. It may be motivated as follows: resource j generates feedback signals indicating congestion and a copy is sent to each user. The latter reacts to such a feedback signal by reducing $x_r[t]$ as mentioned above.

The simulations of the rates of flow and the shadow prices for five random routes and links are shown respectively in Figure 5.1.

When κ is small, there are no oscillations but the system takes a very long time to converge. As κ increases, the number of oscillations also grows with larger amplitudes but shorter convergence times, until it cannot be increased beyond a certain critical point, denoted by κ^*, otherwise the system will explode. The smaller variations for the shadow prices as compared to those for the flow rates with corresponding values of κ can also be noted. This shows that the shadow prices are more robust than the flow rates for the same values of κ. The present simulations indicate that κ may be selected in the interval $(0.005, 0.05)$ for the stability of this model with joint feedback.

As κ increases further, the time taken to converge to the unique stable solution again increases. Moreover, the maximum amplitude which the $x_r(t)$'s reach from zero, before settling down, grows rapidly, until the system no longer converges when $\kappa = 0.1$. In fact, in this particular case, stability is transient. We can also observe that at the round-trip delay of 20, a minimum is reached.

Again, from the right hand side plots of Figure 5.1, for very small gains, the shadow prices for the various resources have the same behaviour as the flow rates, with less

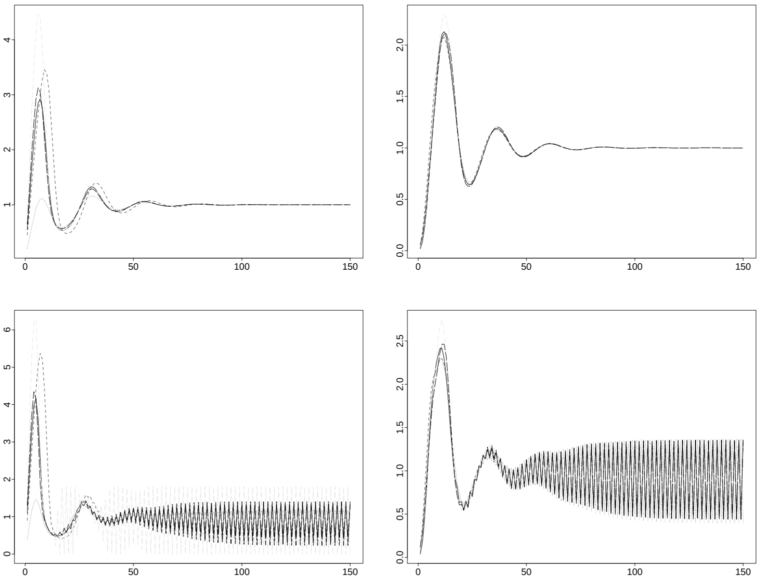

Figure 5.1: Simulation of the rates of flow, x (left hand side graphs), and the shadow prices, μ (right hand side graphs), for κ = 0.05, 0.1, for the deterministic primal algorithm with joint feedback. Note the smaller variations of the shadow prices, as well as the increasing number of oscillations for larger κ. The horizontal axis represents time, while the vertical axis corresponds either to the flow rates, or to the shadow prices.

variability. For most of the links, the shadow prices follow the same paths. On a microscopic scale near convergence, the paths for the various links show staircase behaviours either downwards (if they have values beyond unity) or upwards (if they have values less than unity) until they are exactly unity. The graphs, however, do not reveal much about the relationship between convergence properties and the v's (which measure the number of routes through a resource).

Next, we consider the stochastic version of the primal algorithm considered above. Here, resource j generates feedback signals indicating congestion as a time-dependent Poisson process at rate $\varepsilon_j^{-1}\mu_j[t]$ according to equation (2.29) in Section 2.2.3. When resource j generates a feedback signal, a copy is sent to each user r whose route passes through resource j; and user r reacts to such a feedback signal by reducing $x_r[t]$ by an amount $\kappa\varepsilon_j x_r[t]$.

The importance of the stochastic models is to observe the behaviour of the flow rates and shadow prices about the stable point (in this case unity). The simulations show how the speed of convergence and variability at equilibrium can be controlled. Proper scaling of the simulation parameters can be made to meet realistic dimensions, with similar observations.

The algorithm for the model with stochastic effects is given by

$$x_r[t+1] = x_r[t] + \kappa\Big(w_r - x_r[t]\sum_{j\in J} A_{jr}N_j[t-d(j,r)]\Big) \tag{5.3}$$

where $N_j[t]$, as j and t vary, are independent Poisson random variables with mean $\mu_j[t]$ (note that$N_j[t] = Poisson(\mu_j[t])$), where

$$\mu_j[t] = p_j\left(\sum_{s\in R} A_{js}x_s[t-d(j,r)]\right), \tag{5.4}$$

$w_r = \sum_{j\in r} A_{jr}$ and $p_j(y) = \dfrac{y}{\sum_{s:j\in s} A_{js}}$. The shadow prices are obtained from the same equation (5.2), with $x_s[t] = x_s[0]$ if $t < 0$.

In the experiment for the deterministic model, we have obtained the critical gain parameter as $\kappa^* = 0.01$. For $\kappa = 0.01$, the graphs for the routes and the links are represented as the two lower panels in Figure 5.2. It is obvious from this figure, that there is more variability in the flow rates than in the shadow prices, confirming that the latter are more resistant to fluctuations.

5.1.2 Congestion indication with individual feedback

In our next example, we consider the primal algorithm with individual feedback, discussed in Section 2.2.3. The deterministic discrete time-lag algorithm is given by

$$x_r[t+1] = x_r[t] + \kappa\left(w_r - \sum_{j\in J} A_{jr}x_r[t-2d(j,r)]\mu_j[t-d(j,r)]\right) \tag{5.5}$$

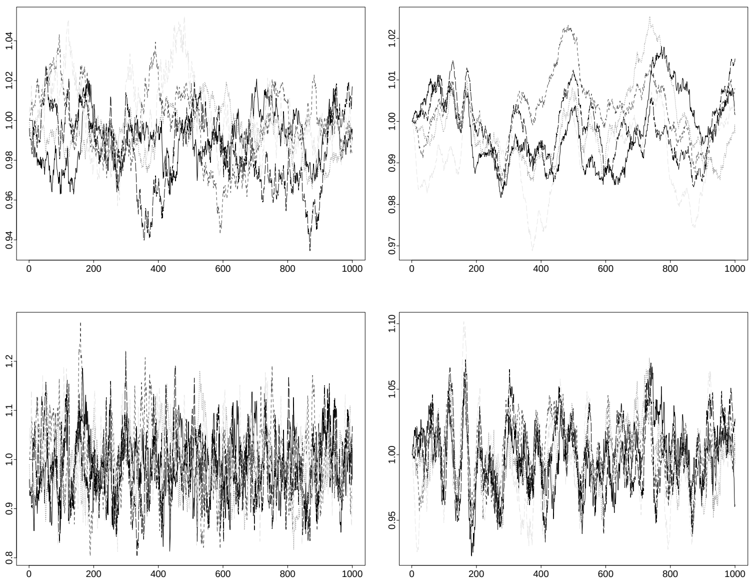

Figure 5.2: Flow rates, x (left.hand side graphs), and shadow prices, μ (right hand side graphs), for the stochastic primal case with joint feedback, $t = 1000$, $\kappa = 0.005, 0.01$. Observe the increase in density and size of amplitudes as κ increases.

and equation (5.2), where $\mu_j[t] = \mu_j[0]$ and $x_r[t] = x_r[0]$ if $t < 0$. Also $d(j,r)$ is the delay between resource j and source r.

In this case, the flow rate increases linearly by an amount κw_r and decreases at a rate proportional to κ, when feedback signals from resource j to user r arise at a rate $x_r[t - 2d(j,r)]\mu_j[t - d(j,r)]$.

Following the same approach as for the joint feedback case, Figure 5.3 shows the variation of the rates of flow (left hand side graphs), and the shadow prices (right hand side graphs) for five random routes and links respectively. As κ varies, some interesting observations can be made.

When κ is small, the convergence times to the stationary solution are very long for both the links and the routes. As in the joint feedback model, we observe that the smaller the w_r, the faster the duration of convergence for the flow rates. This is in accordance to the analytical study in Section 3.3.1. As the curve approaches

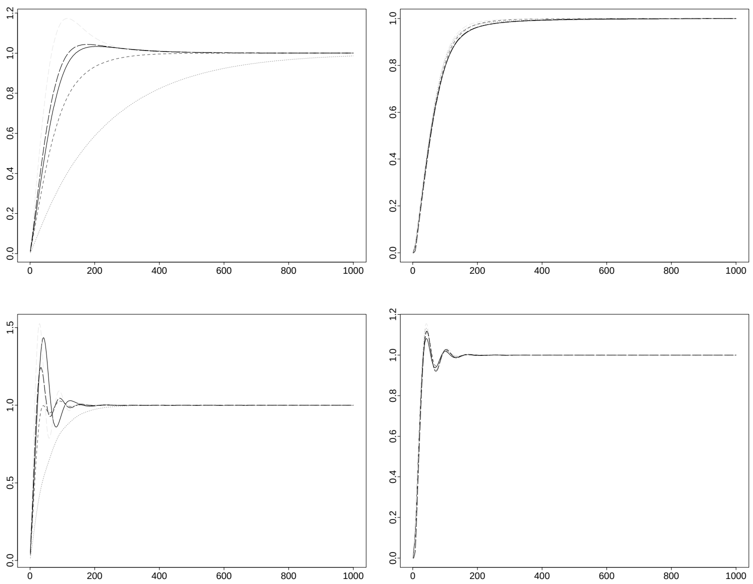

Figure 5.3: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the deterministic primal algorithm with individual feedback, $\kappa = 0.001, 0.004$.

stationary, the convergence follows a step function (of stepsize 0.0001), and the longer it remains in the transition state.

A further increase in κ shows oscillatory behaviour about the stability point. The number and magnitude of oscillations increase with the gain parameter. The effects of the delay included in the algorithm can be observed in the plots for the rates of flow. For some κ in the neighbourhood of κ^*, the system is seen to oscillate symmetrically and periodically about the stable point.

The above model is modified further to observe the effect of noise in the system. In the individual feedback model with stochastic effects, feedback signals from resource j to user r arise at rate $\varepsilon_j^{-1} x_r[t]\mu_j[t]$; and user r reacts to such feedback signal by reducing $x_r[t]$ by an amount $\kappa\varepsilon_j$. This is exactly the system (2.33) in Section 2.2.3.

The algorithm for the stochastic time-lag model is given by

$$x_r[t+1] = x_r[t] + \kappa \left(w_r - \sum_{j \in J} A_{jr} N_{jr}[t - d(j,r)] \right), \tag{5.6}$$

where $N_{jr}[t]$ as j, r and t vary, are independent Poisson random variables with mean $x_r[t - d(j,r)]\mu_j[t]$, and $\mu_j[t]$ is defined by equation (5.2). Moreover,

$$w_r = \sum_{j \in r} A_{jr} \text{ and } p_j(y) = \frac{y}{\sum_{s: j \in s} A_{js}}.$$

The equation used for simulation is given by

$$x_r[t+1] = x_r[t] + \kappa \left(w_r - Poisson \left(\sum_{j \in J} A_{jr} x_r[t - 2d(j,r)] \mu_j[t - d(j,r)] \right) \right), \tag{5.7}$$

with $\mu_j[t] = \mu_j[0]$ if $t < 0$. Figure 5.4 shows the simulations of the rates of flow and the shadow prices for five random routes and links respectively as κ varies.

As κ increases, the number of oscillations along with their amplitudes for the rates of flow also increases until the system is no longer stable for $\kappa^* = 0.01$. Similarly, the range of oscillations for the shadow prices increases, but at a slower rate. These prices are less sensitive to changes in κ as compared to the rates of flow.

If we compare this stochastic model to the corresponding joint feedback case, we observe that for $\kappa = 0.01$, the range of oscillations for the shadow prices are the same. Moreover, the individual feedback model is much more sensitive for the rates of flow as compared to the joint feedback case. The deterministic individual feedback model is found to be stable for a more restricted range of κ as compared to the corresponding joint feedback case. This is explained by the longer delay involved in the individual feedback algorithm (explained by the factor of two involved in equation (5.5)).

There is also more variability for the rates of flow as compared to the shadow prices for the same values of κ. The flow rates overshoot the stationary solution to again

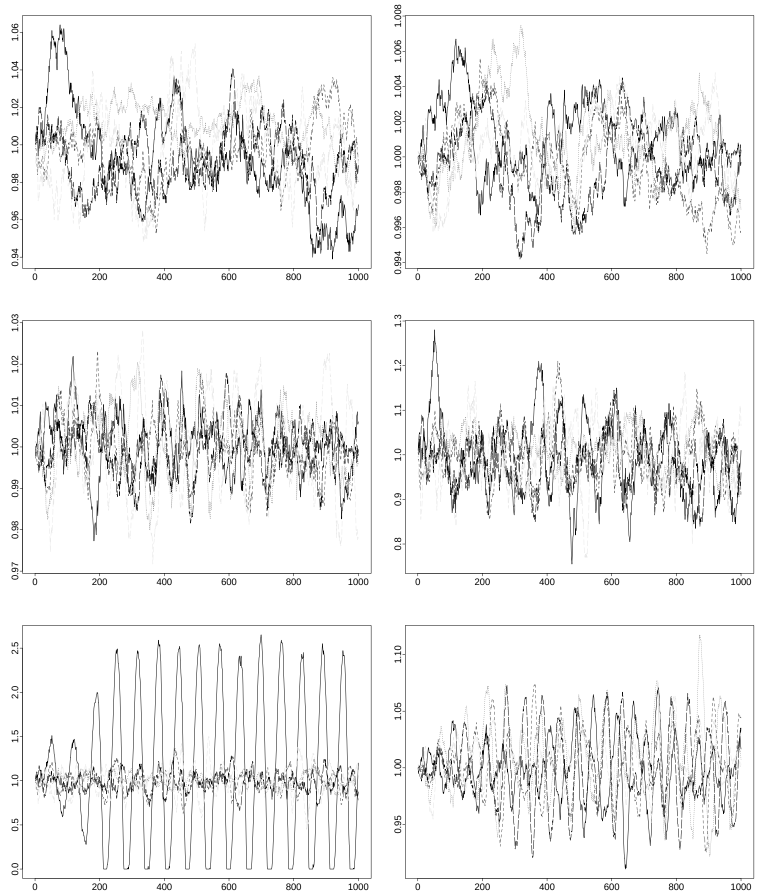

Figure 5.4: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the stochastic primal algorithm with individual Poisson feedback, $\kappa = 0.001, 0.005, 0.01$. Notice the oscillatory behaviour about the stable point. Smaller gain parameters generate fluctuations with smaller amplitude.

converge to stationary in contrast to the shadow prices. In general, the shadow prices are more robust to fluctuations as compared to the flow rates. A comparison of the covariance matrices for the shadow prices and the flow rates may explain this behaviour. Figure 5.5 shows clearly the delay involved in the system for the rates of flow and the shadow prices at equilibrium, although convergence is not reached. A modification of this model, assuming that all the users share the same round-trip delay, leads to a *phase-locked* system, whereby all the rates (or shadow prices) oscillate exactly in phase with each other [37, 60].

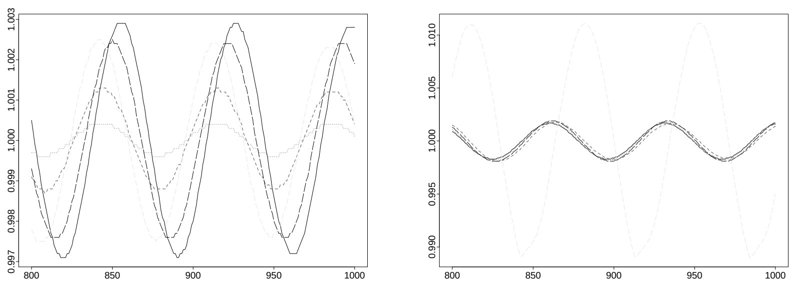

Figure 5.5: Flow rates, x, and shadow prices, μ for the deterministic primal case with individual feedback, $\kappa = 0.005$, showing delay in system. If all connections have the same delay, then we have a phase-locked system.

5.2 A TCP-Like Rate Control Algorithm

Since most existing protocols for the Internet implement variants of the primal algorithm (2.13)–(2.14) (or (2.20)–(2.21)), it is interesting to observe the behaviour of one variant for the random network. There exist tremendous literature on the current protocols used for the Internet, see for example [17, 20, 36, 57, 67, 70].

Assuming the usual notations, the corresponding algorithm is given by

$$x_r[t+1] = x_r[t] + \kappa\big(w_r - \sum_{j\in r} x_r[t-\tau_r]\mu_j[t-\tau_r+d(j,r)]\big) \tag{5.8}$$

where $\mu_j[t] = \mu_j[0]$ if $t < 0$, and

$$\mu_j[t] = \sum_{s:j\in s} x_s[t-d(j,s)]/v_j, \tag{5.9}$$

and $x_r[t] = x_r[0]$ if $t < 0$. The round-trip time for connection r is denoted by τ_r, and is given by

$$\tau_r = 2\, \underset{j\in r}{max}\{d(j,r)\} \;\; \forall\, r \in R, \tag{5.10}$$

where $d(j,r)$ is the delay between resource j and end-node r. In the system of equations (5.8)–(5.10), the flow rate, $x_r[t]$, increases linearly by an amount κw_r and decreases at a rate proportional to κ when feedback signals from resource j to user r arise at a rate $x_r[t-\tau_r]\mu_j[t-\tau_r+d(j,r)]$.

This deterministic system is simulated for different values of κ, over a long period of time, to ensure convergence. A particular value of κ is then chosen as the critical gain parameter from this range. Another gain parameter, denoted by κ_{min} gives the fastest convergence to stability. It also provides a heavy damping and ensures a small number of oscillations with acceptable amplitudes. The parameter can be taken as the optimal gain parameter for the system.

A sample of the simulations for the rates of flow and the shadow prices as κ varies from 0.001 to 0.0035 are shown in Figure 5.6. The simulations in Figure 5.6 show some very interesting features for the end-nodes and the resources of the network respectively.

As κ increases from 0.001, it can be clearly observed that the speed of convergence increases for all the routes at the cost of higher amplitudes and more oscillations. However, this process continues up to a value of $\kappa = 0.0025$. Then, the number of oscillations increases exponentially with higher amplitudes and lighter damping until $\kappa = 0.004$ when the system explodes to never reach stability. It has also been observed that with the round-trip time factor considered, the route length is not directly related to the speed of convergence. There is the combined effect of delay, route length and the round-trip time itself which jointly affect the rate of convergence for any particular route. It means that if we have two routes of the same length but different round-trip times and delays from various resources on the respective routes, one may converge much faster than the other. However, if a route is very long, the amplitude of oscillations are much higher, and this may be combined with a heavy damping effect caused by the choice of κ to equilibrium (in this case unity). This observation has motivated us to study a single bottleneck resource with two end-connections only, in order to understand the behaviour of round-trip time with the

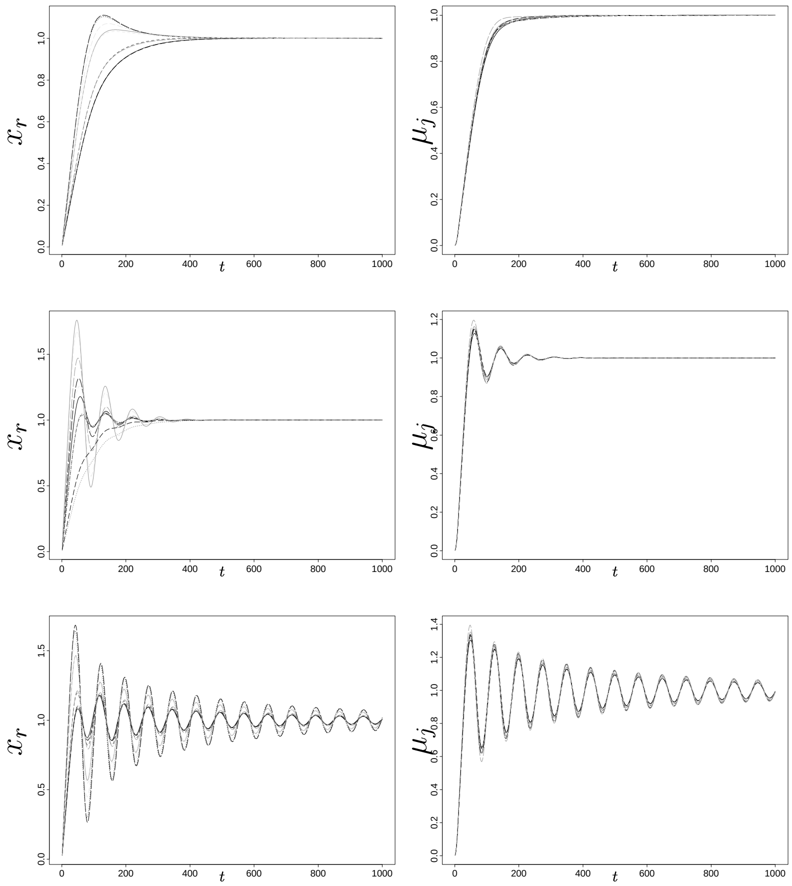

Figure 5.6: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the deterministic primal algorithm for TCP, $\kappa = 0.001$, 0.0025 and 0.0035. Observe the increase in the number and magnitude of oscillations with κ.

gain parameters. This was studied in detail in Chapter 4.

The shadow prices are the main contributors to the feedbacks generated by a resource and sent to a user. In general, it can be observed that the shadow prices are more clustered and oscillate together. The delays are not that obvious. Otherwise, since the shadow prices depend on the rates of flow for particular routes, they are closely related. If observed more closely, the following points may be highlighted. When the gain parameter is very small, the convergence time is very long for all the shadow prices without oscillations. They do not overshoot the stability line as compared to the rates. As the gain parameter increases, the number of oscillations increases.

However, the amplitudes of oscillations are more restrictive for the shadow prices. The growth in the number and size of oscillations continue until $\kappa = 0.005$, when there is no convergence at all in either case. However, both the rates and the shadow prices oscillate periodically with very large amplitudes (of about twice the stable value). Moreover, the shadow prices are not truncated at zero as for the flow rates. For this model, it can be concluded that $\kappa^* = 0.003$. The stochastic model for this algorithm shows similar behaviour as those for the previous models.

It has also been observed from the experiment that the routes with shorter round-trip times have the following characteristics in a chosen time period: less oscillations, lower amplitudes, faster rate of convergence, and more damping.

Hypothesis *For an appropriate value of the gain parameter, a route with smaller round-trip time achieves stability faster with fewer oscillations and more damping.*

This hypothesis also follows from the analytical study in the previous chapter for a single bottleneck resource.

5.2.1 A generalisation of the TCP-like algorithm

In this section, we consider a more general version of the system (5.8)–(5.10) studied in Section 5.2. The model is represented by the following system of differential equations.

$$\frac{d}{dt}x_r(t) = \kappa \left(w_r - \zeta \sum_{j \in J} A_{jr} x_r(t - \tau_r) \mu_j(t - \tau_r + d(j,r)) \right) \tag{5.11}$$

where

$$\mu_j[t] = p_j \left(\sum_{s \in R} A_{js} x_s[t - d(j,r)] \right). \tag{5.12}$$

In this system, ζ is the multiplicative decrease factor. For TCP, ignoring the effect of exponential increase during slow start, and assuming that all packets have the same size, we have $\zeta = 0.5$, [54]. This parameter has the effect of slowing the rate of multiplicative decrease when the flow rates have crossed the stable point. This algorithm gives proportionally fair flows to the connections.

We simulate the discrete version of the above system. The equations representing the system are

$$x_r[t+1] = x_r[t] + \kappa\left(w_r - \zeta \sum_{j\in J} A_{jr} x_r[t-\tau_r]\mu_j[t-\tau_r+d(j,r)]\right) \tag{5.13}$$

where ζ is a positive real constant, $\mu_j[t] = \mu_j[0]$ if $t < 0$,

$$\mu_j[t] = \sum_{s:j\in s} x_s[t-d(j,s)]/v_j,$$

and

$$\tau_r = 2\, \underset{j\in r}{max}\{d(j,r)\} \;\; \forall\, r \in R,$$

with $x_r[t] = x_r[0]$ if $t < 0$.

The above model is investigated in the same usual way for various ζ. Throughout this experiment we assume the maximum delay to be 10. The choice of ζ as a real positive quantity is also justified by the additive increase and multiplicative decrease type of algorithm describing the convergence process for a particular network [57].

The graphs in Figure 5.7 show the relationships between various parameters involved in the process. In general, it is observed that for any particular choice of ζ, both the flow rates and the shadow prices converge to the same equilibrium point $1/\sqrt{\zeta}$ for the gain parameter κ_{min}, which is the gain which gives fastest convergence and highest stability. Let κ^T_{min} and κ^S_{min} be the theoretical and simulated values of κ_{min}, the critical gain parameter, for any particular ζ. From the simulated results, a least-squares best fit for the critical gain parameter was found to satisfy the formula

$$\kappa^T_{min} = \frac{2\times 10^{-3}}{\sqrt{\zeta}}. \tag{5.14}$$

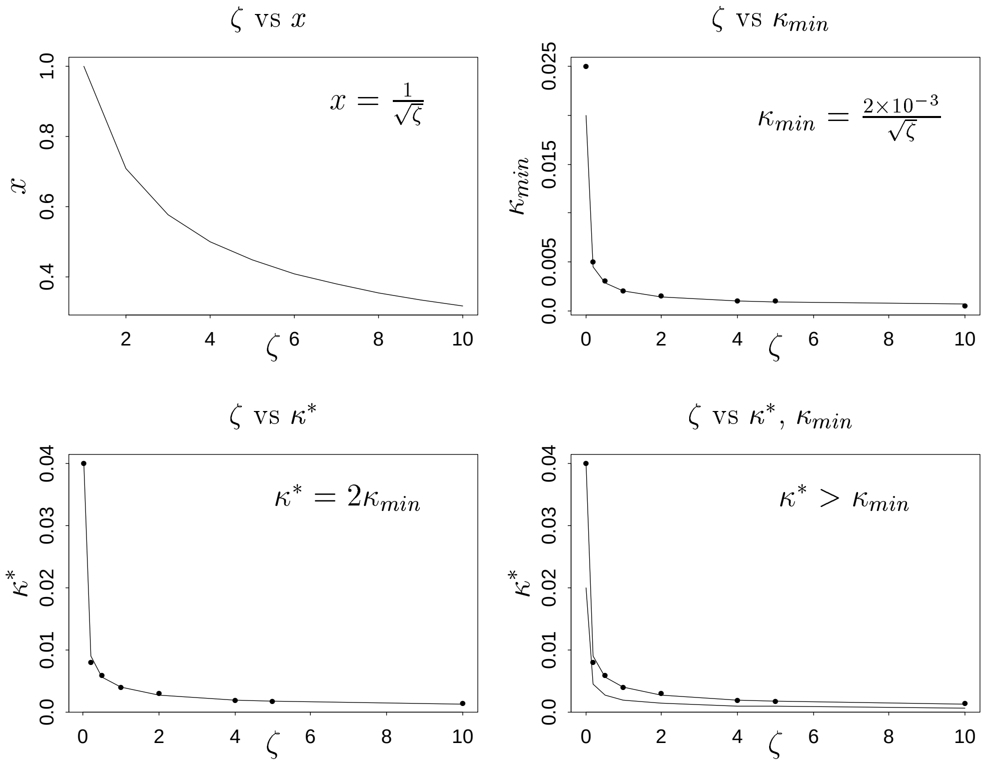

Figure 5.7: Relationship between parameters in the general primal model.

The upper right hand panel in Figure 5.7 illustrates this behaviour. The scatter plot represents the simulated values of κ_{min} while the line represents the best fit equation (5.14). Simulation of the system with different number of routes, $|R|$, suggests that $\kappa_{min} = 2/(|R|\sqrt{\zeta})$.

Assuming that κ^* is the asymptotic (critical) gain parameter of the system for a particular ζ, it is observed that, as in the case of κ_{min}, a best fit for κ^* is $\kappa^*_T = 2\ \kappa^T_{min}$, where κ^*_T is the theoretical value of κ^*. Let κ^*_S be the value of κ^* obtained by simulation. The panel in Figure 5.7 shows how κ^*_S and κ^*_T vary with ζ. As in the previous case, the dots represent simulated values of κ^* while the line represents its best fit curve.

The fourth panel in Figure 5.7 shows the variation of κ_{min} and κ^* with ζ. As expected from the previous discussion, the graph for κ^* is higher than that of κ_{min}.

Following the previous discussions and noting that $x = 1/\sqrt{\zeta}$, we have $\kappa^T_{min} = \dfrac{2}{|R|}\,x$

and $\kappa_T^* = \dfrac{4}{|R|}x$. Therefore we expect a linear parallel variation of κ_{min} and κ^* with x.

Next, we study the variation of ζ and x with the amplitude of oscillations denoted by $a_{\zeta\kappa}$. As κ increases, it is found that for a fixed value ζ, the amplitude increases (even after explosion). Furthermore, for any particular κ, the amplitude decreases with an increase in ζ.

Finally, we study the variation of convergence times with κ for different ζ. The curves in Figure 5.8 show how convergence times T_x for particular values of ζ vary with the gain parameter κ. As ζ increases, $\kappa_{min} \to 0$, and the curves follow a parabolic path with two vertical asymptotes; at the time axis and at $\kappa = \kappa^*$ respectively. Also observe that the mouth of the curve narrows as the curves shift to the left. We have seen in Chapters 3 and 4 that the relation between convergence times and gain parameters is a convex function, with the minimum convergence time reached at κ_{min}. We have also shown that there is a very close relationship between the convergence time and the spectral radius of the time-lag system. All these results suggest a strong relationship among the parameters.

Suppose $T_x = f(\zeta, \kappa)$, and for a fixed ζ, $T_x = f_\zeta(\kappa)$. A challenging research work is to obtain an expression for the function f. This function f should satisfy the conditions obtained in the above experiment: for each ζ, $\kappa_{min} = \frac{2}{|R|\sqrt{\zeta}}$ at which T_x is minimum; as $\kappa \to 0$ for each ζ we should have $T_x \to \infty$ smoothly, and as $\kappa \to \kappa^*$ for each ζ, $T_x \to \infty$. Moreover, the conditions derived in Chapter 4 about the spectral radius of the time-lag system can be combined to have a general expression, in terms of the model parameters. This will allow the users to update their quantities interactively with the information provided by the resources to make efficient use of the network.

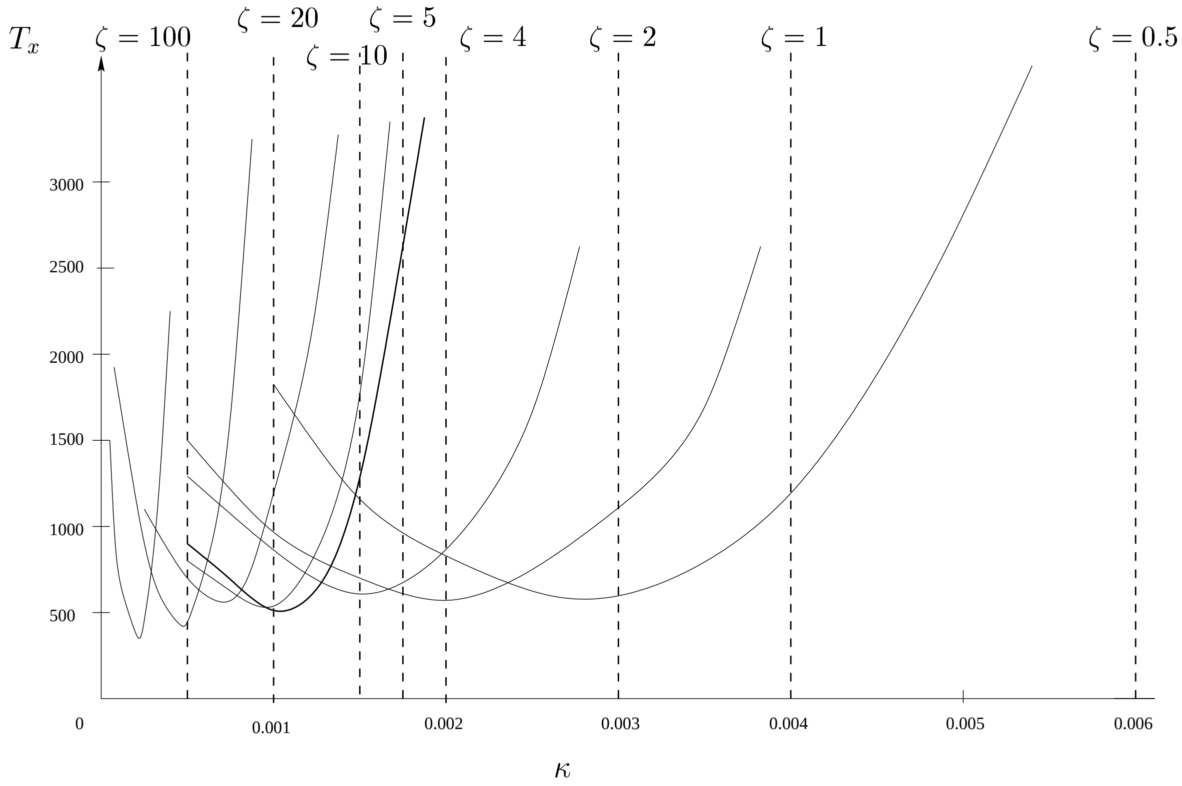

Figure 5.8: Variation of convergence time with κ and $\zeta = 0.5,\ 1,\ 2,\ 4,\ 10,\ 20,\ 100$. The dotted lines represent the critical gains corresponding to the different values of ζ.

5.3 The Dual Algorithm

In Chapter 2, we have derived conditions for the stability and rate of convergence for the system of differential equations representing the dual algorithm (Section 2.3.2). It was deduced that the rate of convergence was proportional to the gain parameter κ. The stochastic versions of this dual algorithm were also studied with the conclusion that the covariance matrix for the shadow prices increases linearly with the gain parameter κ (Section 2.3.3). The time-lag version of the dual algorithm were also analysed in Section 2.3.4. The derived conditions for these specific networks can be applied to any network having these structures.

In this section, we simulate the deterministic time-lag dual algorithm (2.52)–(2.53), for the random network and the regular network described in Chapter 3 respectively. For each type of network, we proceed in two stages: firstly, we simulate the deterministic version of the algorithm and determine the range of the gain parameter over which the system is stable. Subsequently, we simulate the stochastic versions of the algorithms for the gain parameters within the stability interval, to study the behaviour of the flow rates and the shadow prices in the system. For the stochastic version of the algorithm, the feedbacks are generated from Poisson random variables [29].

5.3.1 The dual algorithm for a random network

Consider the discrete time-lag model of the dual algorithm (2.40)–(2.41) for simulation. The following deterministic algorithm, which is the system (2.52)–(2.53), is investigated

$$\mu_j[t+1] = \mu_j[t] + \kappa\Big(\sum_{r\in R} A_{jr}x_r[t-d(j,r)] - \mu_j[t]v_j\Big) \tag{5.15}$$

with $x_r[t] = x_r[0]$ if $t < 0$, where $d(j,r)$ is the delay between link j and route r, $q_j(\mu_j[t]) = \mu_j[t]v_j$ and

$$x_r[t] = \frac{w_r}{\sum_{k\in J} A_{kr}\mu_k[t-d(k,r)]}. \tag{5.16}$$

In this model, the shadow price $\mu_j[t]$ increases at a rate proportional to κ and decreases multiplicatively at a rate $\kappa\mu_j[t]v_j$. As for the primal algorithm, we simulate the deterministic case to obtain the range of the gain parameter κ for stability of the system.

The graphs in Figure 5.9 show the convergence of the rates of flow, x (left hand side graphs), and the shadow prices, μ (right hand side graphs), for κ taking values 0.001, 0.01 and 0.011 for the deterministic dual algorithm defined by equation (5.15).
When κ is very small, the convergence is very slow for both the rates of flow and the shadow prices. As κ increases to 0.01, the number of oscillations increases. This

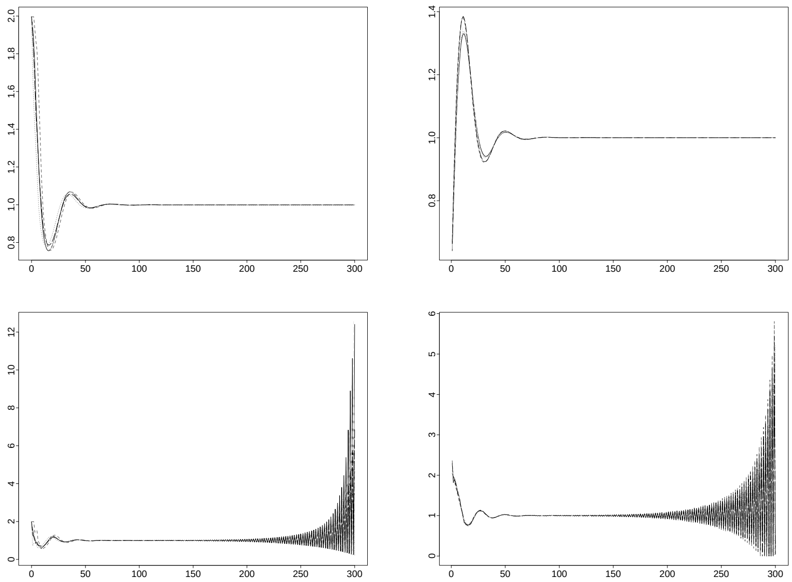

Figure 5.9: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the deterministic dual algorithm for the random network, $\kappa = 0.001$ and 0.011.

behaviour persists until $\kappa = 0.011$, when the system converges for a while to diverge afterwards. For any larger value of κ, the system is unstable. The value of $\kappa^* = 0.01$ is taken as the critical gain parameter. Thus, for the dual system to be stable, a proper choice of κ is in the range $[0.001, 0.01]$. The behaviour is comparable to that of the deterministic primal algorithms for convergence of flow rates and shadow prices for the random network. Also observe that the shadow prices for all the resources follow the same path.

Next, we consider the stochastic version of the dual equation (5.15) for simulation where the feedback is generated by a Poisson random variable. The system is represented by

$$\mu_j[t+1] = \mu_j[t] + \kappa\Big(\sum_{r \in R} A_{jr} N_r[t - d(j,r)] - q_j(\mu_j[t])\Big) \tag{5.17}$$

where $(N_r[t], t = 1, 2, \dots)$, for $r \in R$, are a collection of independent Poisson random

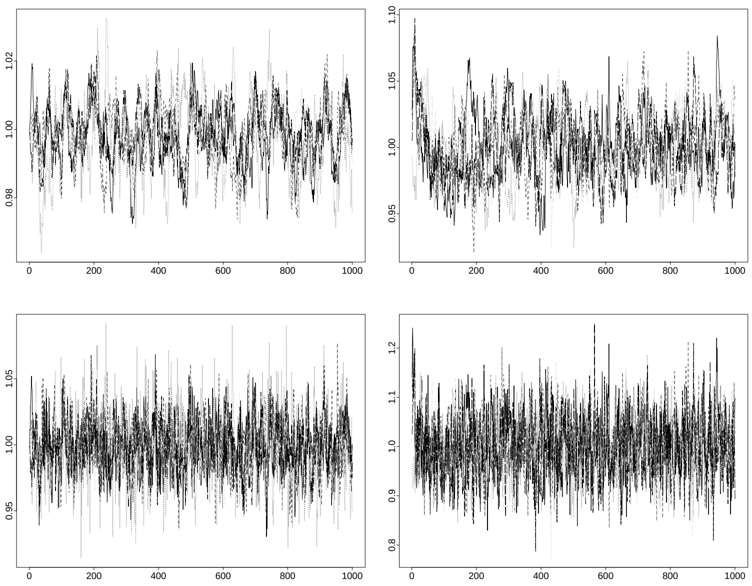

Figure 5.10: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the stochastic dual algorithm for the random network for $\kappa = 0.001$ and 0.005.

variables with mean $x_r[t]$. Note that $x_r[t]$ is again defined by equation (5.16). Figure 5.10 shows the graphs of the rates of flow for five randomly chosen routes and the shadow prices for five random links for κ being 0.001 and 0.005 simulated over time horizons of 1000 time units.

The larger the gain parameter, the larger the number of oscillations about the stable point with higher amplitude, up to $\kappa^* = 0.01$. At this point, the solution diverges from the equilibrium point. Thus, in addition to the deterministic case which suggests a range [0.001,0.01] for κ, this exercise suggests that it is preferable for κ to take values less than 0.01 as shown in the figures, so that the long term oscillations about the stationary solution is small (within some bounds). This helps to maintain stability at equilibrium with acceptable amplitudes of oscillations. Notice the density of the graphs for larger κ. Observe that the simulation of the dual algorithms reveal the smaller variability in flow rates than the shadow prices complementary to the simulation of the primal algorithms.

5.3.2 The dual algorithm for a regular network

The analysis of the random network has shown many features, some of which can be clearly explained while others are still unclear. To study some of the facets which might have been overlooked in the random network case and to confirm some phenomena observed therein, it is appropriate to investigate some networks having definite structures. The structure of the regular network in the present section, and its importance in our analysis were described in Chapter 3. We proceed in the same way as for the random network. The same deterministic and stochastic algorithms (5.15)–(5.16) are used.

The above system is simulated for a regular network with 10 nodes. The upper two panels in Figure 5.11 show the convergence of the rates of flow, x (left hand side graph), and the shadow prices, μ (right hand side graph), for $\kappa = 0.01$. The lower two panels represent the simulation of the stochastic system for ten randomly chosen routes and links respectively, for κ being 0.01.

For the deterministic system, we observe that the shorter the routes (the lowest curve of each graph in the figure), the shorter the time for convergence. In fact, the length of the routes vary from 1 to $|J|$. The nth curve numbered $r + (n-1)|J|$, $r = (1, 2, \dots, |J|-1)$, $n = (1, 2, \dots |J|)$ overlap with the rth curve, so that only $|J|$ curves appear in the figures for the rates of flows. As κ increases, the time for convergence decreases sharply at the expense of more oscillations with higher amplitudes to finally diverge for $\kappa \geq 0.15$.

For this particular network, the number of routes passing through a particular node is fixed and is given by $\sum_{l=1}^{|J|} l$ ($= 55$ when $|J| = 10$). All the curves follow the same path. Further, the convergence times for the shadow prices are much less than for the corresponding curves for the rates of flows for particular values of κ. The convergence times follow the same behaviour for both the flow rates and the shadow prices, and achieve a minimum for both when $\kappa_{min} = 0.1$. For this network, $\kappa^* = 0.15$ is critical as the system is unstable beyond this point.

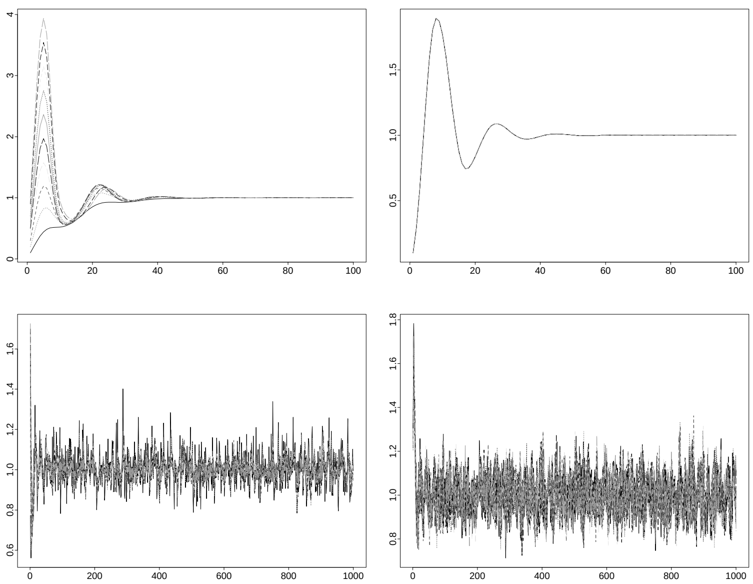

Figure 5.11: Flow rates, x (left hand side graphs), and shadow prices, μ (right hand side graphs), for the dual algorithm for the regular network for $\kappa = 0.1$.

The size of the oscillations for the rates of flow is about the same as for shadow prices. However, the curves for the shadow prices are more dense to that of the rates of flow, possibly because of the delays involved in the system.

The simulation and analysis of two different types of networks for the dual algorithm have revealed some very interesting features. The theoretical results are supported for both types of networks. Even the observations which appeared awkward for the random network, such as the behaviour of κ^* with a near the origin, have been confirmed by the regular network. These features will need further study. These results may be very important to help in congestion control for large communication networks such as the Internet and ATM networks.

The simulations of the deterministic algorithms determine the range for the gain parameter κ while the stochastic cases identify the long term stability of the networks. As κ increases, the density as well as the magnitude of oscillations increase. Moreover, the flow rates and shadow prices behave similarly (symmetry about the stable point) for the stochastic models.

Another interesting aspect of the study is that there is a convex relationship between the gain parameters and the convergence times. Thus, the simulations for extremely small κ will have a very long convergence time, but without oscillations. As the gain parameter approaches its critical value, the convergence time again increases, but this time with lots of variability.

The experimental support in this chapter has shown the importance of the many aspects of congestion control studied in this thesis. The experiments performed for the primal algorithms can also be performed for the dual algorithms, and vice-versa, with similar results. With proper scaling of the model parameters, a relatively well controlled system can be obtained which will be useful in the designing of feedback control systems for congestion avoidance in large-scale communication networks. The simulations confirm that the dynamical systems are well-controlled for both the random and the regular network.

Chapter 6

Conclusion

In this dissertation we have addressed the issue of how the available bandwidth within a large-scale broadband network should be shared between competing streams of elastic traffic. An optimisation framework leads to a decomposition of the overall system problem into a separate problem for each user, in which the user chooses a charge per unit time that he is willing to pay, and one for the network; we have shown that two classes of rate control algorithm are naturally associated with the objective functions appearing in, respectively, the primal and the dual formulation of the network's problem.

The main result obtained concerns the stability and fairness of the two classes of rate control algorithms. Those simple linear increase/multiplicative decrease schemes are stable about a system optimum characterised by a proportional fairness criterion: they may be interpreted in terms of either congestion indication feedback signals or explicit rates based on shadow prices. We have shown that stability and convergence of these algorithms may be established using well-defined Lyapunov functions which are essentially the network's objective function. We have also studied the stability of algorithms in the presence of stochastic perturbations and time-lags.

The algorithms developed in Chapter 2 can be generalised to large-scale communication networks. Due to the size and scope at which networks such as the Internet are expanding, there is a need to design protocols which can be easily implemented in future networks. For this purpose, we were motivated to analyse our rate control algorithms for random networks and compare the results with those of a regular network.

The analysis has shown that the algorithms are stable under properly chosen contract parameters, assuming that users are adaptive. Simple results about the eigenvalues of the linearised systems in terms of the model parameters indicate how the overall system can be monitored by simple changes in the variables. However, the analysis becomes very cumbersome when the gain parameter is no longer fixed for all the users, and no relationships can be deduced from the simulations of the time-lag systems. This has motivated the investigation of a single bottleneck resource to understand some microscopic details which may have been overlooked.

The results in Chapter 4 are based on a TCP-like rate control algorithm, where we assume that the model parameters are chosen by each user, to make efficient use of the network. We use the concept of packet marking as congestion indicators. Several scenarios of the models were studied for two users by varying the gain parameters and the round-trip times. The conclusions are surprising but very helpful for the end-users. They indicate how the various routes may choose the gain parameters, if the round-trip times of the selected routes are known, so as to maintain stability while using a fair share of the bandwidth.

The simulation results presented in Chapter 5 have illustrated the stability of variants of the primal and the dual algorithms respectively. They show how small changes in the parameters do affect stability as well as the rate of convergence. In particular, we have seen how the gain parameters affect the convergence rates and how small variations in each algorithm can influence the behaviour of the system. The gain parameters were assumed fixed in Chapter 3. However, when the gain parameters are route dependent, the analysis becomes more complicated and simulations of the random matrix do not reveal any relationship with the delays or round-trip times. In Chapter 4, we have studied this particular aspect by considering a single active bottleneck resource. By assuming the gain parameters to be fixed, the system becomes simpler to analyse. However, we cannot control the level by which the flow rates or shadow prices overshoot the stable point. For this reason, we have introduced a damping factor which controls the rate at which the steady point is crossed. This modified model has been studied thoroughly by simulation and the importance of this new parameter revealed.

To summarise, we can maintain that the results obtained provide important insights into the behaviour of the system, and allow us to quantify the various relationships and trade-offs between the rate of convergence of the algorithms, the magnitude of fluctuations about the stable point, and the stability of the system.

The work in this thesis has shown many features of the algorithms which need further exploration and understanding. Interesting and challenging questions remain concerning the stability of the entire system under more general assumptions on users' reactions to the rates allocated to them by the network, and when the number of users and the amounts of capacity available for elastic traffic vary randomly. A long term issue is the implementation of the results on an actual network such as the Internet. A body of work is emerging in this direction: see Crowcroft and Oeschlin [28], Gibbens and Kelly [44, 45], Key and McAuley [69] and Key et al. [70]. Also, there is the enabling work of Ramakrishnan and Floyd [87] which proposes adding Explicit Congestion Notification (ECN) to the Internet protocol. The aim is to make the end-users choose the parameters in order to make efficient use of the network.

Our investigation is far from conclusive and has raised several issues for further research. An interesting research problem, for example, is to determine how we can choose κ_r so as to minimise the spectral radius of the matrix L. Another challenging question is the generalisation of the work in Chapter 4 for $|R|$ end-users, and ultimately, for a network with an arbitrary number of resources and routes.

Bibliography

[1] L. Arnold. On Wigner's Semicircle Law for the eigenvalues of random matrices. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, **19**:191–198, 1971.

[2] L. Arnold. Deterministic version of Wigner's Semicircle Law for the distribution of matrix eigenvalues. *Linear Algebra and its Applications*, **13**:185–199, 1976.

[3] K.J. Arrow and F.H. Hahn. *General Competitive Analysis.* Oliver and Boyd and Holden Day, Edinburgh, 1971.

[4] S. Athuraliya, D. Lapsley, and S. Low. An enhanced random early marking algorithm for Internet flow control. Available from `http://www.ee.mu.oz.au/staff/slow/research/projects.html`, 2000. To appear in INFOCOM 2000.

[5] S. Athuraliya and S. Low. Optimization flow control with Newton-like algorithm. Available from `http://www.ee.mu.oz.au/staff/slow/`, 1999. To appear in Telecommunication Systems.

[6] S. Athuraliya and S. Low. Random early marking. Available from `http://www.ee.mu.oz.au/staff/slow/research/projects.html`, 1999.

[7] Z.D. Bai, J.W. Silverstein, and Y.Q. Yin. A note on the largest eigenvalue of a large dimensional sample covariance matrix. *Journal of Multivariate Analysis*, **26**(2):166–168, August 1988.

[8] Z.D. Bai and Y.Q. Yin. Limiting behaviour of the norm of products of random matrices and two problems of Geman-Hwang. *Probability Theory and Related Fields*, **73**:555–569, 1986.

[9] Z.D. Bai, Y.Q. Yin, and P.R. Krishnaiah. On limiting spectral distribution of product of two random matrices when the underlying distribution is isotropic. *Journal of Multivariate Analysis*, **19**:189–200, 1986.

[10] S. Bajaj, L. Breslau, and D. Estrin. Improving simulation for network research. Technical Report 99-702, USC Computer Science Department, March 1999.

[11] S. Barnett. *Matrices: Methods and Applications.* Oxford University Press, 1990.

[12] A. Berman, M. Neumann, and R. Stern. *Non-negative Matrices in Dynamic Systems.* John Wiley and Sons, 1989.

[13] D. Bertsekas and R. Gallager. *Data Networks.* Prentice-Hall, Englewood Cliffs, New Jersey, 1987.

[14] P. Billingsley. *Convergence of Probability Measures.* John Wiley and Sons, 1968.

[15] J.-C. Bolot and A.U. Shankar. Dynamic behavior of rate-based flow control mechanisms. *ACM Computer Communications Review*, **20**:35–49, 1990.

[16] F. Bonomi, D. Mitra, and J.B. Seery. Adaptive algorithms for feedback-based flow control in high-speed wide-area networks. *IEEE Journal of Selected Areas in Communications*, **13**:1267–1283, 1995.

[17] L.S. Brakmo and L.L. Peterson. TCP Vegas: End-to-end congestion avoidance on a global Internet. *IEEE Journal of Selected Areas in Communications*, **13**(8):1465–1481, October 1995.

[18] M. Carmeli. *Statistical Theory and Random Matrices.* Marcel Dekker, 1983.

[19] A. Charny. An algorithm for rate allocation in a packet-switching network with feedback. Master's thesis, Massachusetts Institute of Technology, USA, May 1994.

[20] A. Charny, D.D. Clark, and R. Jain. Congestion control with explicit rate indication. In *Proceedings ICC 95*, pages 1954–1963, June 1995. Available from `http://www.cis.ohio-state.edu/~jain/papers/charny.htm`.

[21] A. Charny, K.K. Ramakrishnan, and A. Lauck. Time scale analysis and scalability issues for explicit rate allocation in ATM networks. *IEEE/ACM Transactions on Networking*, **4**(4):569–581, August 1996.

[22] D-M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, **17**:1–14, 1989.

[23] S. Chong, R. Nagarajan, and Y-T. Wang. First-order rate-based flow control with dynamic queue threshold for high-speed wide-area ATM networks. *Computer Networks and ISDN Systems*, **29**(18):2201–2212, February 1998.

[24] D. Clark. Adding service discrimination to the Internet. *Telecommunications Policy*, **20**, 1996. A version is available from `http://ana-www.lcs.mit.edu/anaweb/abstracts/TPRC2-0.html`.

[25] C. Courcoubetis, G. Kesidis, A. Ridder, J. Walrand, and R.R. Weber. Admission control and routing in ATM networks using inferences from measured buffer occupancy. *IEEE Transactions on Communications*, **43**:1778–1784, 1995.

[26] C. Courcoubetis, V.A. Siris, and G.D. Stamoulis. Integration of pricing and flow control for available bit rate services in ATM networks. In *Proceedings IEEE Globecom '96*, London, UK, 1996.

[27] C. Courcoubetis and R.R. Weber. Buffer overflow asymptotics for a switch handling many traffic sources. *Journal of Applied Probability*, **33**:886–903, 1996.

[28] J. Crowcroft and Ph. Oechslin. Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP. *ACM Computer Communications Review*, **28**:53–67, 1998.

[29] L. Devroye. *Non-Uniform Random Variate Generation.* Springer-Verlag, New York, 1986.

[30] K. Fall and S. Floyd. Simulation-based comparisons of Tahoe, Reno, and SACK TCP. *Computer Communications Review*, **26**(3):5–21, July 1996.

[31] C. Fang and P.R. Krishnaiah. Asymptotic distributions of functions of the eigenvalues of some random matrices for non-normal populations. *Journal of Multivariate Analysis*, **12**:39–63, 1982.

[32] K.W. Fendick, M.A. Rodrigues, and A. Weiss. Analysis of a rate-based feedback control strategy for long-haul data transport. *Performance Evaluation*, **16**:67–84, 1992.

[33] S. Floyd. TCP and Explicit Congestion Notification. *ACM Computer Communications Review*, **24**(5):10–23, October 1994.

[34] S. Floyd. Internet research: Comments on formulating the problem. With contributions from D. Estrin and G. Minshall, January 1998.

[35] S. Floyd. Congestion control principles. Internet Engineering Task Force, Internet Draft, January 2000. A version of the draft is available from `http://www.aciri.org/floyd/`.

[36] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking*, **7**(4):458–472, August 1999.

[37] S. Floyd and V. Jacobson. On traffic phase effects in packet-switched gateways. *Internetworking: Research and Experience*, **3**(3):115–156, 1992.

[38] S. Floyd and V. Jacobson. Random Early Detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, **1**(4):397–413, August 1993.

[39] Z. Gajic and M.T.J. Qureshi. *Lyapunov Matrix Equation in System Stability and Control*, volume **195** of *Mathematics in Science and Engineering*. Academic Press, 1995.

[40] S. Geman. A limit theorem for the norm of random matrices. *The Annals of Probability*, **8**:252–261, 1980.

[41] S. Geman. Almost-sure stable oscillations in a large system of randomly coupled equations. *SIAM Journal of Applied Mathematics*, **42**(4):695–703, 1982.

[42] S. Geman. The spectral radius of large random matrices. *The Annals of Probability*, **14**(4):1318–1328, 1986.

[43] S. Geman and C.R. Hwang. A chaos hypothesis for some large systems of random equations. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, **60**:291–314, 1982.

[44] R.J. Gibbens and F.P. Kelly. Distributed connection acceptance control for a connectionless network. In P.B. Key and D.G. Mason, editors, *Proceedings of 16th International Teletraffic Congress*, pages 941–952. Elsevier, 1999. Available from `http://www.statslab.cam.ac.uk/~frank/dcac.html`.

[45] R.J. Gibbens and F.P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, **35**:1969–1985, 1999.

[46] R.J. Gibbens, R.A. Mason, and R. Steinberg. Multiproduct competition between congestible networks. Discussion paper in Economics and Econometrics 9816, University of Southampton, September 1998. Also available from `http://www.soton.ac.uk/~ram2/index.html`.

[47] P. Glendinning. *Stability, Instability and Chaos: An Introduction to the Theory of Non-linear Differential Equations.* Cambridge University Press, 1994.

[48] S.J. Golestani and S. Bhattacharya. A class of end-to-end congestion control algorithms for the Internet. In *Proc. Sixth International Conference on Network Protocols.* Available from `http://www.bell-labs.com/user/golestani/`, 1998.

[49] U. Grenander and J.W. Silverstein. Spectral analysis of networks with random topologies. *SIAM Journal of Applied Mathematics*, **32**:499–519, 1977.

[50] F. Hahn. Stability. In K.J. Arrow and M.D. Intriligator, editors, *Handbook of Mathematical Economics*, volume **II**, pages 745–793. North-Holland, Amsterdam, 1982.

[51] J. Hale. *Theory of Functional Differential Equations.* Springer-Verlag, New York, 1977.

[52] E.J. Hernandez-Valencia, L. Benmohamed, S. Chong, and R. Nagarajan. Rate control algorithms for the ATM ABR service. *European Transactions on Telecommunications*, **8**(1):7–20, January/February 1997.

[53] M.W. Hirsch and S. Smale. *Differential Equations, Dynamical Systems and Linear Algebra.* Academic Press, New York, 1974.

[54] P. Hurley, J-Y. Le Boudec, and P. Thiran. A note on the fairness of additive increase and multiplicative decrease. Available from `http://icawww.epfl.ch/ICApublic/asp/publication`, 1999.

[55] C-R. Hwang. A brief survey on the spectral radius and the spectral distribution of large random matrices with i.i.d. entries. In *Random Matrices and Their Applications*, volume **50** of *Contemporary Mathematics*, pages 145–152. American Mathematical Society, 1984.

[56] International Telecommunications Union, Geneva. *Traffic Control and Congestion Control in B–ISBN*, 1996. Recommendation 1371.

[57] V. Jacobson. Congestion avoidance and control. *Computer Communications Review*, **18**(4):314–329, August 1988.

[58] R. Jain. Congestion control and traffic management in ATM networks: Recent advances and a survey. *Computer Networks and ISDN Systems*, **28**(13):1723–1738, October 1996. A version of this paper is available from `http://www.cis.ohio-state.edu/~jain/papers.cnis.htm`.

[59] R. Jain, K.K. Ramakrishnan, and D-M. Chiu. Congestion avoidance in computer networks with a connectionless network layer. Technical Report DEC–TR–506, Digital Equipment Corporation, USA, June 1997.

[60] R. Johari and D. Tan. End-to-end congestion control for the Internet: Delays and stability. Preprint.

[61] D. Jonnson. Some limit theorems for the eigenvalues of a sample covariance matrix. *Journal of Multivariate Analysis*, **12**:1–38, 1982.

[62] D. W. Jordan and P. Smith. *Nonlinear Ordinary Differential Equation*. Clarendon Press, Oxford, 2nd edition, 1977.

[63] F.P. Kelly. On tariffs, policing and admission control for multiservice networks. *Operations Research Letters*, **15**:1–9, 1994.

[64] F.P. Kelly. Modelling communication networks, present and future. *Philosophical Transactions of the Royal Society London A*, **354**:437–463, 1996.

[65] F.P. Kelly. Notes on effective bandwidths. In S. Zachary F.P. Kelly and I. Ziedins, editors, *Stochastic Networks: Theory and Applications*, volume **4** of *Royal Statistical Society Lecture Notes*, pages 141–168. Oxford University Press, Oxford, 1996.

[66] F.P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, **8**(1):33–37, January/February 1997.

[67] F.P. Kelly. Models for a self-managed Internet. Available from `http://www.statslab.cam.ac.uk/~frank/smi.html`, December 1999. Royal Society, London.

[68] F.P. Kelly, A.K. Maulloo, and D.K.H. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, **49**:237–252, 1998.

[69] P.B. Key and D. McAuley. Differential QoS and pricing in networks: where flow-control meets game-theory. *IEE Proceedings Software*, **146**(1):39–43, February 1999.

[70] P.B Key, D. McAuley, P. Barham, and K. Laevens. Congestion pricing for congestion avoidance. Technical Report MSR-TR-99-15, Microsoft Research, February 1999. Available from `http://research.microsoft.com/pubs/`.

[71] D. Lapsley and S. Low. Random Early Marking for Internet congestion control. Available from `http://www.ee.mu.oz.au/staff/slow/research/projects.html`, 1999. Preprint.

[72] S.H. Low. Optimization based congestion control for the Internet. *IEEE Network*, August 1999. To appear. A version of the paper is available from `http://www.ee.mu.oz.au/staff/slow/research/projects.html`.

[73] S.H. Low. Optimization flow control with on-line measurement or multiple paths. In P.B. Key and D.G. Mason, editors, *Proceedings of 16th International Teletraffic Congress*. Elsevier, 1999. Also available from `http://www.ee.mu.oz.au/staff/slow/`.

[74] S.H Low and D.E Lapsley. Optimization flow control, I: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, **7**(6):861–874, December 1999.

[75] H. Lütkepohl. *Handbook of Matrices.* John Wiley and Sons, 1996.

[76] J.K. MacKie-Mason and H.R. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet.* Prentice-Hall, Englewood Cliffs, New Jersey, 1994. Also available from `http://www-personal-umich.edu/~jmm/papers.html`.

[77] L. Massoulié and J. Roberts. Bandwidth sharing: Objectives and algorithms. Available from `http://www.res.enst.fr/~oueslati/Pub/elastoc.ps`, 1999. To appear in INFOCOM'99.

[78] R. Mazumdar, L.G. Mason, and C. Douligeris. Fairness in network optimal flow control: optimality of product forms. *IEEE Transactions on Communications*, **39**(5):775–782, May 1991.

[79] L.W. McKnight and J.P. Bailey. *Internet Economics.* MIT Press, Cambridge, Massachussetts, 1997.

[80] M.L. Mehta. *Random Matrices.* Academic Press Limited, London, UK, revised and enlarged 2nd edition, 1991.

[81] R.K. Miller and A.N. Michel. *Ordinary Differential Equations.* Academic Press, New York, 1982.

[82] J. Mo, R. La, V. Anantharam, and J. Walrand. Analysis and comparison of TCP Reno and Vegas. Available from `http://www.path.berkeley.edu/~jhmo`, July 1998. To appear in INFOCOM'99.

[83] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. A version of the paper is available from `http://www.path.berkeley.edu/~jhmo`, April 1998. Preprint.

[84] A. Odlyzko. A modest proposal for preventing Internet congestion. Available from `http://www.research.att.com/~amo/doc/complete.html`, September 1997. AT & T Laboratory Research.

[85] A.M. Odlyzko. Paris Metro Pricing for the Internet. In *Procedings ACM Conference on Electronic Commerce (EC'99)*, pages 140–147. ACM, 1999. Also available from `http://www.research.att.com/~amo/doc/paris.metro.pricing.ps`.

[86] V. Paxson and S. Floyd. Why we don't know how to simulate the Internet. In *Proceedings of the 1997 Winter Conference*. Atlanta, GA, 1997.

[87] K.K. Ramakrishnan and S. Floyd. A proposal to add Explicit Congestion Notification (ECN) to IP. Request For Comments RFC 2481, January 1999.

[88] J. Roberts and L. Massoulié. Bandwidth sharing and admission control for elastic traffic. Available from `http://www.res.enst.fr/~oueslati/Pub/itcjap.ps`, 1998.

[89] S. Shenker. Fundamental design issues for the future Internet. *IEEE Journal of Selected Areas in Communications*, **13**:1176–1188, 1995.

[90] J.W. Silverstein. On the eigenvectors of large dimensional sample covariance matrices. *Journal of Multivariate Analysis*, **30**(1):1–16, July 1989.

[91] J.W. Silverstein. On the weak limit of the largest eigenvalue of a large dimensional sample covariance matrix. *Journal of Multivariate Analysis*, **30**(2):307–311, August 1989.

[92] G. Stépán. *Retarded dynamical systems: stability and characteristic functions*. Number 210 in Pitman Research Notes in Mathematics Series. Longman Scientific and Technical, London, UK, 1989.

[93] D.K.H. Tan. *Mathematical Models of Rate Control for Communication Networks*. PhD thesis, St John's College, University of Cambridge, UK, August 1999.

[94] H.R. Varian. *Microeconomic Analysis*. Norton, New York, 3rd edition, 1992.

[95] K.W. Wachter. The strong limits of random matrix spectra for sample matrices of independent elements. *The Annals of Probability*, **6**(1):1–18, 1978.

[96] R. Wegmann. The asymptotic eigenvalue distribution of a certain class of random matrices. *Journal of Mathematical Analysis and Applications*, **56**:113–132, 1976.

[97] E.W. Weistein. *CRC Concise Encyclopedia of Mathematics*. CRC Press, 1999.

[98] P. Whittle. *Optimization Under Constraints*. John Wiley and Sons, Chichester, 1971.

[99] S. Wiggins. *Introduction to Applied Non-linear Dynamical Systems and Chaos*. Springer-Verlag, New York, 1977.

[100] E.P. Wigner. Characteristic vectors of bordered matrices with infinite dimensions. *Annals of Mathematics*, **62**:548–564, 1955.

[101] D. Wischik. How to mark fairly. Available from `http://www.statslab.cam.ac.uk/~djw1005`, 1999.

[102] E. Wong and B. Hajek. *Stochastic Processes in Engineering Systems*. Springer-Verlag, 1985.

[103] Y.Q. Yin. Limiting spectral distribution for a class of random matrices. *Journal of Multivariate Analysis*, **20**:50–68, 1986.

[104] Y.Q. Yin, Z.D. Bai, and P.R. Krishnaiah. Limiting behaviour of the eigenvalues of a multivariate F matrix. *Journal of Multivariate Analysis*, **13**:508–516, 1983.

[105] Y.Q. Yin, Z.D. Bai, and P.R. Krishnaiah. On the limit of the largest eigenvalue of the large dimensional sample covariance matrix. *Probability Theory and Related Fields*, **78**(4):509–521, August 1988.

[106] Y.Q. Yin and P.R. Krishnaiah. A limit theorem for the eigenvalues of the product of two random matrices. *Journal of Multivariate Analysis*, **13**:489–507, 1983.