

ABSTRACT

NINAN, BOBBY M. Resource Pricing for Connection-Oriented Networks. (Under the direction of Associate Professor Michael Devetsikiotis).

Network pricing has important implications in the revenue generation, resource management, system optimization and congestion control of computer networks. We depart from the prevalent idea of marginal cost pricing and provide a holistic, bi-level optimization framework to model the interaction between network entities in a connection oriented network. Users are treated as utility maximizing entities who allocate the available bandwidth among themselves by playing a distributed, noncooperative rate game. The ensuing Nash equilibrium is analyzed for the single link Erlang network and the multi-link product form networks. Variants based on the upper bound of the blocking are also studied owing to their role in reducing computational complexity. Theoretical results are then validated using numerical simulation for varying network scenarios. An extension of the rate adaptation game based on Recursive Least Squares is proposed for dealing with the imperfect information scenario. These exhibited favorable convergence, accuracy and scalability properties. Gradient-free schemes are then developed for revenue maximization. These are based on novel stochastic approximation techniques such as Finite Difference Stochastic Approximation (FDSA) and Simultaneous Perturbation Stochastic Approximation (SPSA). It is observed that the network employed price discrimination for optimizing its objective function and partitioning its available capacity among competing users.

Resource Pricing for Connection-Oriented Networks

by

Bobby M. Ninan

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial satisfaction of the
requirements for the Degree of
Doctor of Philosophy

Operations Research program

Raleigh

2004

Approved By:

Dr. Arne Nilsson

Dr. Yannis Viniotis

Dr. Michael Devetsikiotis
Chair of Advisory Committee

Dr. Stephen Campbell

To my parents, who made me who I am.

എന്നെ ഞാനാക്കി തീർത്ത എന്റെ മാതാപിതാക്കൾക്ക്.

Biography

Bobby Muthootu Ninan was born in Kerala, India on March 27, 1979. After studying at CCPLM and Chinmaya Vidyalaya, he passed his ICSE and ISC exams from Udyogamandal school, Eloor. Udyogamandal was a turning point in his life which made him realize the need for a multifaceted personality and a sound technical education. He then entered the portals of Indian Institute of Technology (IIT), Kharagpur for his Bachelors of Technology degree. A fascination for applied mathematics and computers guided him to pursue a PhD degree in Operations Research and Electrical and Computer Engineering. He looks forward to a future where he can continue to pursue interesting research in multidisciplinary areas such as biomathematics, genomics and complex systems.

Acknowledgements

First and foremost, I would like to acknowledge the faith my parents had placed in me while spotting my potential and nurturing it all the way from kindergarten to graduate school. My only regret is that instead of being a doctor of medicine as per my mother's wishes, I ended up being a doctor of operations research and computer networks. May the twain meet some time in future.

My teachers ranging from Udyogamandal school to North Carolina State University played a significant part in moulding my values and outlook towards life. Thanks to Abraham Thomas for his impeccable teaching skills which rightly emphasized the fundamentals over rote learning. Vanaja Pillai with her indulgence towards my eccentricities has been a constant source of support and love. IIT introduced me to Dr. A. K. Datta whose dedication towards his students and research still remain unsurpassed. When I was a stranger in a strange land, Dr. Henry Nuttle helped me to find my feet and soothe my fears. During the four years of my PhD, my advisor Dr. Mihail Devetsikiotis has suffered through my mood swings and widely fluctuating interests in disparate fields. Stephen Campbell introduced me to the mathemagical world of control and has stood by me during my darkest hour.

My doctorate would have been impossible if my friends and room mates had not stepped in to preserve my sanity. My gratitude to Yoganand Saripalli, Sai Oruganti and Devender Chauhan knows no bounds. NCSU also introduced me to two people I ended up admiring - Arin Chaudhuri and Srikant Nalatwad. Debating with Arin has always been an intellectual roller coaster, be it the strategic vision for India or the continuum hypothesis. Srikant with his moral fortitude has stood by me through my good times and bad times and taught me that goodness and decency are still worth fighting for.

Last but not the least, a special thanks to Marhn Fullmer for his invaluable help and guidance at a critical juncture in my academic life. He taught me that a doctorate in computer networks will always be incomplete without hands on experience and getting one's hands dirty.

Contents

| | |
|--|------------|
| List of Figures | vii |
| 1 Introduction | 1 |
| 1.1 The Evolution of Network Pricing | 2 |
| 1.1.1 The QoS Paradigm versus Overprovisioning | 3 |
| 1.1.2 Flat versus Usage Pricing | 4 |
| 1.2 The Case for Differential Pricing | 6 |
| 1.3 Related Work | 8 |
| 1.4 Contributions of This Dissertation | 11 |
| 1.5 Outline | 12 |
| 2 Problem Statement | 15 |
| 2.1 Model Formulation | 16 |
| 3 Bandwidth Sharing in Circuit Switched Networks | 19 |
| 3.1 Erlang System | 20 |
| 3.2 Speeding up Computations | 24 |
| 3.2.1 Farago Bound | 26 |
| 3.2.2 Martinez Bound | 29 |
| 3.3 Product Form Networks | 31 |
| 3.4 Results and Discussion | 34 |
| 3.4.1 Convergence of Arrival Rates | 35 |
| 3.4.2 Effect of Bounds on Nash Equilibrium | 38 |
| 3.4.3 Triangle Network | 38 |
| 4 Imperfect Information Regime | 44 |
| 4.1 Erlang System | 45 |
| 4.2 Feedback Based Rate Control | 46 |
| 4.3 Dynamic Estimation | 48 |
| 4.4 Rate Adaptation Algorithms | 48 |
| 4.4.1 Original Algorithm | 48 |
| 4.4.2 Logarithmic Variant | 49 |

| | | |
|----------|---|-----------|
| 4.4.3 | Newton-Raphson Variant | 50 |
| 4.4.4 | Simulation Results | 51 |
| 4.4.5 | Scalability | 51 |
| 4.4.6 | Accuracy | 52 |
| 4.4.7 | Effect of Aggressiveness on System Convergence | 53 |
| 4.4.8 | Impact of User Demand on System Convergence | 54 |
| 5 | Bi-level Optimization | 57 |
| 5.1 | Rate Optimization | 57 |
| 5.2 | Price Optimization | 61 |
| 6 | Stochastic Approximation | 64 |
| 6.1 | Convergence of Stochastic Approximation | 65 |
| 6.1.1 | “Statistics” conditions | 66 |
| 6.1.2 | “Engineering” conditions | 67 |
| 6.2 | Gradient-Free Algorithms | 68 |
| 6.2.1 | Finite Difference SA (FDSA) | 68 |
| 6.2.2 | Simultaneous Perturbation SA (SPSA) | 70 |
| 6.3 | SPSA Pricing algorithm | 71 |
| 6.4 | Numerical Results | 72 |
| 6.4.1 | Efficiency of SPSA | 73 |
| 6.4.2 | Impact of Number of Users on Revenue Generation | 74 |
| 6.4.3 | Effect of Initial Price on Convergence | 75 |
| 6.4.4 | Triangle Network | 76 |
| 7 | Conclusion | 80 |
| 7.1 | Summary | 80 |
| 7.2 | Future Work | 81 |
| | Bibliography | 83 |

List of Figures

| | | |
|------|--|----|
| 1.1 | Access Network. | 13 |
| 1.2 | Network pipe with circuits. | 14 |
| 2.1 | Network-User system interaction. | 16 |
| 3.1 | Roots for Erlang B system. | 25 |
| 3.2 | Graphical illustration of Lemma 2. | 26 |
| 3.3 | Roots of Farago system. | 28 |
| 3.4 | Comparison of bounds with blocking probability. | 29 |
| 3.5 | Comparison of termination criteria. | 35 |
| 3.6 | Arrival rates under resource surplus in a single link | 36 |
| 3.7 | Arrival rates under resource deficit in a single link | 36 |
| 3.8 | Convergence to NEP. | 37 |
| 3.9 | Speed of convergence to NEP. | 38 |
| 3.10 | Convergence to equilibrium rates. | 39 |
| 3.11 | Triangle network. | 39 |
| 3.12 | Arrival rates under resource surplus in a triangle network. | 41 |
| 3.13 | Bandwidth allocation under resource surplus in a triangle network. | 42 |
| 3.14 | Arrival rates under resource deficit in a triangle network. | 42 |
| 3.15 | Bandwidth allocation under resource deficit in a triangle network. | 43 |
| 4.1 | Feedback based user rate adaptation. | 46 |
| 4.2 | Comparison of measured and estimated outputs for 2 users. | 49 |
| 4.3 | Iterations vs Number of users: Logarithmic and Newton-Raphson. | 52 |
| 4.4 | Error vs Number of users: Logarithmic and Newton-Raphson variants. | 53 |
| 4.5 | Iterations vs β , Original algorithm. | 54 |
| 4.6 | Iterations vs β , Logarithmic variant. | 55 |
| 4.7 | Iterations vs β , Newton-Raphson variant. | 55 |
| 4.8 | Iterations vs ϵ , Logarithmic variant. | 56 |
| 5.1 | Bandwidth Supply and Demand | 59 |
| 5.2 | Mapping of Θ^* to feasible bandwidth region | 60 |
| 5.3 | Network-User interaction | 63 |

| | | |
|-----|--|----|
| 6.1 | Comparison of efficiency for FDSA and SPSA. | 73 |
| 6.2 | Revenue generation for varying number of users using FDSA. | 74 |
| 6.3 | Revenue generation for varying number of users using SPSA. | 75 |
| 6.4 | Market clearing effects for a system of 10 users. | 76 |
| 6.5 | Effect of Initial Price on Convergence. | 77 |
| 6.6 | Capacity allocation in the Bottleneck Link. | 78 |
| 6.7 | Price and Revenue Variation in the Bottleneck Scenario. | 78 |
| 6.8 | Sensitivity of bandwidth allocation to maximal rates in the bottleneck link. | 79 |

Chapter 1

Introduction

The meteoric rise of Internet companies and their concomitant meltdown have brought renewed attention to the concept of network pricing. It has been advanced as a possible panacea for the ills plaguing the networking sector ranging from revenue generation, congestion control to enabling Quality of Service criteria for network traffic. The success of microeconomic policies in controlling a noisy, distributed system like the global markets quite akin to the Internet bears testimony to this notion. Efforts are ongoing to hasten and render seamless the transition from a simple, flat pricing scheme to a socially efficient, usage based regime.

The motivation behind the pricing of computer networks is closely interlinked with their history, service architecture and future expansion. Although connection oriented networks¹ had their beginning in the public telephone network, the communication arena has been revolutionized by the advent of the Internet. It has popularized the notion of packet switching and spawned a number of technical and commercial innovations. Its pervasive influence has also managed to alter the field of connection oriented networks which in their latest incarnation supplement and coexist with the Internet. Thus any decision to radically alter the makeup and philosophy of the “Net” will also have serious implications to connection oriented networks and vice versa.

In this chapter, we provide a brief introduction of the evolution of pricing and its relation with Quality of Service and the ensuing debate between flat and usage based

¹In this dissertation, they are alternatively referred to as “circuit switched networks” or “loss networks”.

pricing. We argue that price differentiation is the key to asset recovery and profits in the beleaguered telecommunications sector. Related work in the academic literature is reviewed for illustrating the development of pricing algorithms for network resource management. We then enumerate the contributions of this dissertation towards this objective and conclude with a road map for the rest of this work.

1.1 The Evolution of Network Pricing

The Public Switched Telephone Network (PSTN), the precursor of the current Internet, was built and regulated by monopolies all over the world. PSTN is an example of a circuit switched network where a physical path is dedicated for transmission between two end points for the duration of the transmission. It offered a single type of service namely telephony which in turn led to a uniform network architecture. Since nationwide telephony networks were operated either by government agencies or private monopolies, they ensured a homogeneity in service and the possibility of centralized control. This in turn led to a simplified, uniform tariff structure and billing policy. It also ushered in a regime where intelligence in the form of circuit reservation and routing was concentrated in the network while the endpoints (telephone sets) were relegated to the role of passive players.

In contrast, the connectionless datagram principle was developed mainly to ensure network reliability, a key concern for Advanced Research Project Agency Network's (ARPANET) survivability in the face of a nuclear war. But it was also driven by the evolving economics of transmission costs and switching devices. When faced with the prospect of cheap transmission lines relative to switches, the telephone industry came up with connection oriented networks where a large number of lines interfaced with few switches to create end-to-end circuits. However as routers became inexpensive and bandwidth prices soared it made more sense to increase utilization by means of statistical multiplexing. This led to the development of packet switched networks like the present Internet.

The Internet has come a long way from its humble beginning of a defense controlled research network. Today it spans the globe and has made a successful transition into a vibrant social and commercial infrastructure. This spectacular success can be attributed to a combination of open standards, interoperable architectures and the "end-to-end" principle. For the sake of scalability, the complexity of the network has now been pushed to the edges

leading to a “dumb” network populated with intelligent end points.

Government funding of the Internet came to an end when on December 23rd 1992, the National Science Foundation expressed its intention to stop supporting the ANS T3 backbone in the near future. Telecom companies like Sprint and AT&T soon jumped into the infrastructure bandwagon hoping to garner a slice of the Internet backbone pie. Several business plans based on future earnings were proposed and lapped up in the exuberant investment climate predating the ‘bubble’. The projected demand never materialized and when the bubble broke, the companies were straddled with large quantities of dark fibre. Surprisingly as of 2003, the \$80 billion revenues from wireless far outstrip the \$35 billion from Internet [36, 37]. This illustrates one of the principal reasons for Internet pricing namely the imperative for drawing up a feasible *revenue generation* plan for mitigating the cross subsidization of data traffic by voice traffic.

1.1.1 The QoS Paradigm versus Overprovisioning

While telecom companies were driven by the objectives of price discrimination and thereby return on investments, academia’s interest was piqued by the relevance of pricing as a tool for network control. The fledgling Internet experienced its first severe congestion in 1987 prompting the development of a congestion control algorithm as part of the Transmission Control Protocol (TCP) suite. It was realized that such problems would be exacerbated by the rise of applications with ever increasing appetite for bandwidth. The severe delay problems in NSFNET during November 1992 due to some audio/video broadcasts served to illustrate these concerns. The need of the hour was a mechanism designed to encourage a socially optimal solution wherein high value bits (e.g., telemedicine packets carrying life saving information) would be given preference over others. The field of providing Quality of Service (QoS) in the Internet was thus born.

Taking a cue from PSTN, the Internet was designed as a best effort system with the network not providing any guarantees on the timeliness or even the arrival of packets. The QoS paradigm however required a network that could carry out service differentiation with packets serviced depending upon their value. But incentives were necessary to prevent users from inflating their packet values and requesting for better service. Price discrimination of services was found to be ideal for encouraging service differentiation with the associated revenues paying for any needed network expansions. The prevalence of inelastic

applications like interactive audio/video necessitated the introduction of admission control schemes akin to those deployed in telephony. The need for extending the service model with users explicitly requesting service is detailed in [42]. It was also suggested that the basic best-effort architecture be left intact with QoS schemes solely reserved for resource intensive high quality real-time services.

Since the notion of differentiated services demanded changes to the prevailing network architecture, a section of the networking community offered overprovisioning as a possible panacea for congestion. Bandwidth was becoming increasingly cheap due to economies of size as more and more users were joining the Internet. Further the advent of novel optical technologies like DWDM could squeeze more and more bandwidth into the same fibre. Under the assumption of “almost free” bandwidth, it was believed that huge overprovisioning would be economically feasible. The startling implications of measurements from the BellCore network [25] pointed to the high variability and possible self-similarity of data traffic. This burstiness thus indicated that any overprovisioning of capacity based on peak characteristics would be far costlier than the usual average based allocation.

1.1.2 Flat versus Usage Pricing

As the Internet was a public good, the academic community tried to follow the footsteps of economists by resolving to maximize the social welfare of its users. Congestion was seen as the playing out of the classic “tragedy of commons” where individual users with unrestricted access overgrazed the system to the detriment of others. This could be alleviated by a usage based scheme with users getting charged for the amount of traffic they consume. For maintaining social optimality these charges would have to be set equal to the marginal cost of usage. Since bandwidth scarcity occurs only during congestion, this marginal cost is essentially the same as the congestion cost.

The notion of congestion pricing was developed to account for the social costs imposed by the user on the rest of the population during periods of congestion. Several usage based schemes [10] were introduced to promote social optimality. One of the schemes which caught attention was the Vickrey auction based ‘smart market’ developed by MacKie-Mason and Varian [27]. This required that each user indicate the value of her packets by incorporating a bid in the packet header. The routers would then allow all packets whose bid exceeded the marginal cost to enter the network. This marginal cost would be equal

to the congestion cost imposed by the next arriving packet. The router would however charge all the admitted packets only the marginal cost maintaining optimality. Users have no incentive to under-report their bids as admission on the network depends on an unknown and possibly higher price. Packets which were rejected could wait for transmission at a less congested period thereby trading dollars for delay.

On the other end of the spectrum was the idea of flat pricing with users enjoying unlimited access after paying an access fee. This scheme though suboptimal was conceptually simpler to usage based pricing as it was compatible with the existing architecture obviating the need for extensive monitoring and accounting mechanisms. It is argued that usage based schemes ran counter to the risk aversion and need for predictability of consumers. Odlyzko cites examples involving several networks like mail, telegraph and telephone services [35], railroad and highways [38] to argue for the inexorable march towards simplicity. When prices are kept simple and low, more and more users migrate leading to profits from increased revenues. The increase in user population increases the value of the network as expounded by Metcalfe's law thus leading to a positive spiral.

A critique of the optimality paradigm pervading the pricing literature was provided by Shenker in [43]. Most of the Internet backbone is owned by profit maximizing companies with little interest in socially optimal schemes. Since most of the costs for maintaining the infrastructure consisted of fixed costs, it is not clear whether marginal costs would be able to recover the operating costs. The utility derived by users from individual packets depend on the delay faced by them, a variable inherently difficult to predict. If they are part of a flow, their individual utilities would be influenced by the delivery of the rest of the flow. The inaccessibility of marginal cost severely curtail the implementation of schemes like the 'smart market' forcing researchers to look for alternative schemes. Any optimal pricing mechanism would need to be deployed globally, an idea which would require extensive standardization and runs counter to the idea of Internet being a collection of heterogeneous networks. Edge pricing proposes to reduce complexity by shifting the mechanisms to the edge. Monitoring and billing policies are simplified by employing a scheme based on expected congestion (like time-of-day pricing) and expected path. It rejects the perceived dichotomy of usage and flat pricing by considering them as competing design choices for pricing at the edge.

1.2 The Case for Differential Pricing

An economic system is considered *Pareto efficient* if the allocation of resources has reached such a point where there is no way to make an entity better off without making some one else worse off. Pareto efficiency is achieved by maximizing the *social surplus* which is defined as the sum of the total utilities of the users and provider minus costs. There could be more than one Pareto efficient outcomes as discussed below.

The prevalent theme in network pricing has been economically efficient schemes obtained by setting the price equal to that of the marginal cost in delivering the service. Here we define marginal cost as the cost of providing an incremental unit of the good. The rationale for marginal cost pricing revolved around the idea of a provider who was able to deliver as many units as long as it was profitable for him to do so. Under perfect competition, only those companies who would be able to sell their goods with razor thin profit margins would succeed. This would enable the customers to obtain services at the cheapest possible price. The outcome is thus Pareto efficient with the users getting all the social surplus and the provider making zero profit. Such a business model makes sense only when the sunken charges are lower in comparison to the marginal cost of providing the service. This paradigm may not be feasible in the telecommunications and networking industry [47] since they involve large fixed costs for setting up the infrastructure, economies of scale and near zero incremental cost in servicing an additional customer. Thus setting efficient prices would lead to negligible service charges for offsetting the variable costs incurred. However, this would not be sufficient for paying for the billions of dollars sunk in setting up the communications network in the first place.

Differential pricing is a natural charging scheme with the network charging users based on their willingness to pay. Price differentiation can be classified into:

- **First-degree price discrimination** often referred to as perfect price discrimination, where the producer varies her selling price depending upon the user and lot size. Each unit of good is sold to the user who values it most highly and at the maximum price she is willing to pay for it. Although this would ensure that the producer ends up getting all the surplus, it requires her to know the willingness-to-pay of her customers. This coupled with the ability to prevent resale of goods prevent the adoption of this stratagem in the real world.

- **Second-degree price discrimination** where the producer varies her selling price depending upon the lot size only. Thus all users buying the same amount of good pays the same price. Volume discounts fall into this category. The rationale behind this approach is the potential difficulty to distinguish between customers with differing willingness-to-pay. Instead users are presented with multiple price-quantity packages of differing quality levels providing an incentive for them to self-select. Quality variation could also be a direct result of the pricing policy as evident in Paris Metro Pricing [34].
- **Third-degree price discrimination** where the producer varies her selling price depending upon the user only. For any given user, every unit of the good is sold at the same price. This is the most prevalent form in the telecommunication industry as exemplified by lifeline pricing and differential pricing for individuals and business. Acknowledging its ubiquity, in this dissertation, we employ third-degree price discrimination to maximize network revenue.

The INDEX study [9] gave a fillip to pricing research by providing empirical proof that users were willing to pay based on usage for genuinely better (and strictly guaranteed) Quality of Service under certain circumstances. The advantages of differential pricing are three fold. First, the resultant surplus garnered by the service provider can be used to pay for sunken costs and future expansion. It is hence ubiquitous in industries which possess high fixed costs such as airlines, publishing and telecommunications. Second, is its ability to expand the customer base using cross subsidies. Frequently technological advances require a large amount of investment in research and development. Companies try to recover their costs by heavily charging the early adopters for exclusive access. Later as the products mature, the prices are lowered to include small businesses and individuals. Similarly when telecom companies are required to provide mandatory services for low-income households, they could either be provided with government subsidies or allowed to carry out price discrimination. Subsidies and handouts are often inefficient and encourage corruption. Providers could then avoid losses by cross-subsidizing emergency with appropriate charges on businesses and high-income households. Third, it creates incentives for the users to regulate their service requirements by trading off performance to monetary benefits. In a shared resource network like the Internet, end users should cooperate and adjust transmission rates to avoid congestion collapse. Although the development of Transmission Control

Protocol (TCP) friendly rate adaptation protocols [39] has received considerable interest, there currently exists no incentive for users to deploy them. Further, the advent of Linux and its myriad variants has given more freedom for end users to override the defaults and manipulate their hosts to be less cooperative with their peers. In such a scenario, differential pricing remains one of the very few options to shape user behavior.

1.3 Related Work

The pioneering work of Kelly [18] laid the foundation for analyzing the role of bandwidth pricing in network optimization. A communication network was visualized as a system striving to maximize the utilities of its users while working within its capacity constraints. It was demonstrated that this could be decomposed into several user utility maximization problems mediated by a price per unit bandwidth. The notion of social fairness was popularized by introducing the concepts of max-min and proportional fairness. A feasible allocation vector is termed max-min fair if any increase individual user rates are permitted only if they do not adversely affect those of others. Hence this idea is akin to the idea of Nash equilibrium where an absolute priority is provided to all users, big and small.

The fairness criterion was incorporated in the design of rate control strategies for ATM Available Bit Rate service by Hernandez-Valencia et al. [14]. Proportional fairness on the other hand favors smaller flows less aggressively by permitting modifications to the rate vector whose sum of proportional changes in nonpositive. This criterion was employed by the network to allocate bandwidth to users who explicitly declared their desired bandwidth price or “willingness to pay”. It was shown that solution was proportionally fair if the user utilities were logarithmic functions of bandwidth. In [19], this model was applied in conjunction with Additive Increase/ Multiplicative Decrease (AIMD) rate control algorithms similar to the one employed by the TCP in the current Internet. The network’s objective function was shown to be a Lyapunov function to the dynamic system composed of such AIMD algorithms thereby demonstrating the stability of the system optima. The primal and dual versions of the problem were analyzed and feedback based control strategies developed for their solution. System optimization was carried out by the primal algorithm modifying the rates and the dual algorithm adjusting the resource prices.

Modifications to the existing TCP/IP were proposed in [13] to provide incentives

for users to modify their transmission rates. While the current TCP versions deduce congestion from round trip times and packet drops, the authors favored the use of Explicit Congestion Notifications (ECN) to users. Such implementations of ECN in routers are becoming prevalent after its introduction by Floyd in [12]. Overloaded links convey their state information by marking and charging packets passing flowing through them. The packets are marked in proportion to the marginal cost incurred on the network. This is the probability that the removal of a random packet would reduce the number of packet losses in the link by one. The performance of this scheme was studied for elastic, intermittent users and file transfers.

Low et al. in [26] offered another take on congestion pricing with users modifying their transmission rates in the face of varying bandwidth prices. A decentralized, dual algorithm based on gradient projection was developed enabling individual links to independently set their link prices. Convergence results were established and an asynchronous variant was proposed to counter the effects of time varying feedback delays. Signalling overheads resulting from network-user interactions were minimized in [17] with the network conveying only the number of congested links on the user's path. Thus a single byte of data would be adequate in communicating the system information and that too only in the presence of congestion. When there is no congestion, the users would increase their rates according to the derivative of their utility function. On receiving the congestion notification from the network, it would decrease its sending rate proportional to the number of congested links. Since a subgradient based optimization method was proposed to maximize total user utility, utilities could encompass the wider ambit of continuous functions.

Economists have traditionally employed game theory to analyze the the behavior of users in markets regulated by supply and demand. The users are modeled as rational agents striving to maximize their individual utility functions while competing/cooperating with their counterparts. Such assumptions appear much more reasonable for a collection of computing machines interacting with each other through dedicated communication channels as in the case of the Internet. Cocchi et al. [5] was among the first to bring network pricing under the purview of noncooperative game theory. Our work has been influenced by their game based model of user dynamics and emphasis on service class sensitive pricing titled "priority pricing". However the emphasis was on promoting user efficiency without taking into account the relationship network pricing and user elasticity to price.

A charge sensitive TCP [24] was investigated by La et al. as a likely candidate for

a rate control strategy compatible with the present Internet. It does not require any explicit feedbacks from the network and hence can be incrementally deployed on the Internet by modifying only the end hosts. The network decides on a bandwidth price reflecting the queuing delay imposed by each user on others. In the framework of Kelly, users modify their willingness to pay for a given network price. Since packet delay and hence network price is in turn a function of the user rates, this leads to a noncooperative game between the users. The existence of the subsequent Nash equilibrium was proved for a single bottleneck.

Recently this work was generalized in [2] with users implicitly maximizing the net benefit accrued from data transmission and the subsequent queuing delay faced. Using an idealized fluid model of the network and queues, the global stability of the algorithm in a general network was proved under delay.

Stackelberg games are those in which a leader or supervisor (typically the network) participates in the game typically to steer the users (followers) to some desired behavior. Stackelberg games for communication networks were used in [7,8] where the game's leader was itself a selfish user of network. In [22], the authors consider a Stackelberg game in which users choose routes in a wired network after the leader has chosen routes for its own traffic; in choosing, the leader controls user behavior to optimize some network utility or to achieve some other "global" goal. Their problem formulation admits a closed-form expression for the equilibrium points which can then be steered by the leader to a preordained operating point (incentive compatibility).

In [41], the authors modeled wireless customers by formulating a CDMA power data-rate control games for which the equilibrium point was studied. The game theoretic framework of resource allocation was formalized by Yaïche et al. [50] and extended to include the idea of Nash Bargaining Solutions (NBS) from cooperative game theory. Assumptions on maximizing the aggregate user utilities and resulting proportionally fair solutions were shown to be a special case of the NBS model. Since the utility information is often always not available, the authors proposed an algorithm based on the users' maximum and minimum rates as well as the allocated bandwidth. A dual-based distributed algorithm was then shown to achieve fair bandwidth allocation between users. Proofs of convergence was shown to hold even when the utility functions were not second order differentiable (\mathcal{C}^2).

1.4 Contributions of This Dissertation

A significant portion of the literature has been devoted to the application of pricing to networks running the Internet Protocol (IP). Thus most of the schemes seek to influence transmission rates by modifying the TCP window size. Paralleling the ascendancy of resource pricing has been the growing awareness of the shortcomings of the current Internet in combating congestion and supporting differentiated services and novel traffic streams. Multi Protocol Label Switching (MPLS) [49] and its optical analogue Generalized Multi Protocol Lambda Switching (GMP λ S) have been proposed to be the foundation of the network of the future. These protocols use labels to set up Label Switched Paths which are similar to virtual circuits in a circuit switched network. A MPLS-enabled router examines only the label of the packet while forwarding it. This reduced processing delay is crucial for next generation networks where Gigabits of data are handled every second. Connection oriented networks are thus back in vogue with the realization that any next generation network would need to be a hybrid of packet and circuit switched networks. Integration of a feasible pricing strategy into such an architecture entails modeling the benefits accrued and possible repercussions on user behavior and network stability.

This dissertation investigates the interplay between the two components of a connection oriented system – the service provider referred to as the “network” and the collection of users utilizing the system. The contribution can be summarized as follows:

- In the framework of Cocchi et al., we developed a user model based on noncooperative game theory and third degree price discrimination. User requests for bandwidth were mapped onto the set of call request rates.
- Investigated the corresponding Nash Equilibrium Point of the user game for a single link and a multi-link network using Erlang’s blocking formula and Erlang’s fixed point approximation respectively. The convergence to the fixed point and the impact of maximal rates on the equilibrium were also studied. Theoretical conditions for its uniqueness were derived for the Erlang network.
- Looked into upper bounds such as the ones proposed by Farago and Martinez for reducing the computational complexity for a single link network. The bifurcation point for the Farago system was calculated for finding the operating region of interest.

- Extended the Erlang game to include the imperfect information regime where measurements are corrupted by noise and delay. Three algorithms based on Recursive Least Squares were proposed and their convergence, accuracy and scalability was validated by simulation results.
- Introduced a bi-level optimization model to describe the interaction between the network and its users. Conditions were developed for the maximization of total user utility, its distributed variant and the network revenue maximization problem.
- Two gradient-free algorithms based on the novel idea of stochastic approximation were advanced for solving the network optimization problem. The superiority of the Simultaneous Perturbation variant was demonstrated vis-a-vis its Finite Difference counterpart. Numerical simulations of the network-user complex illustrated revenue maximization and price differentiation.

In addition to the above mentioned MPLS/GMPAS combine, other possible applications of this work include Virtual Private Networks overlaid over the public Internet, Wireless/PCS networks. The advent of broadband Internet services over Digital Subscriber Lines (DSL) and cable has sparked interest in access networks such as in Fig 1.1. For this single bottleneck link, Fig. 1.2 illustrates the logical circuits of multiple users.

1.5 Outline

The rest of the thesis is organized as follows:

In the next chapter, we introduce the formal problem statement of the network-user interaction. This includes an abstract version of the user rate adaptation game. Chapter 3 studies this formulation in the setting of bandwidth sharing in circuit-switched networks. The single link Erlang system, its upper bound variants and multi-link counterpart are then treated in detail. We then present our foray into the imperfect information scenario in Chapter 4 where adaptive algorithms modify their rates in the presence of noisy measurements. The bi-level optimization model unifying the rate and price optimization paradigm is formulated in Chapter 5. Stochastic Approximation schemes for revenue maximization are presented in Chapter 6. Finally in Chapter 7, we present our conclusion and suggest avenues for future work.




Figure 1.1: Access Network.

The following notation is used throughout this thesis. Vectors are represented in upper case. For example, if Θ is a vector, $\Theta(k)$ represents its k^{th} iterate, θ_k its k^{th} component. Both optimal (Θ^*) and equilibrium (Θ^{eqm}) values are represented by superscripts.




Figure 1.2: Network pipe with circuits.

Chapter 2

Problem Statement

Traditionally the dimensioning of connection-oriented networks has been implemented after taking into account the long term demand, available capacity and the required QoS as per the service level agreements (SLA) with its customers. This arrangement was then coupled with a Connection Admission and Control (CAC) mechanism which regulated the incident traffic into the network. One of the major reasons for the scalability of the current Internet has been its reliance on a dumb network populated with intelligent endpoints. This approach differs from that of connection-oriented networks like the Public Switched Telephone Network (PSTN) wherein the intelligence in terms of CAC and routing are concentrated in the network.

Our model envisages a network wherein noncooperative users allocate resources among themselves in the absence of a centralized CAC scheme. Knowledge of the utility function of a user is considered private and is not exchanged with other users or the network. The performance metric of any user would thus be a function of those of other users as well¹. We envisage resource allocation as the result of a two stage process:

- Price Optimization: The network chooses a price to maximize its revenue, the product of total bandwidth consumed and market price.
- Rate Optimization: For a given fixed network charge $\$M \equiv [M_1, \dots, M_{|R|}]$, the users

¹Again, we emphasize that we use the term “users” in an abstract way to denote any entity, physical or virtual, individual or aggregate, that accesses the network resources as one unit from the point of view of pricing and QoS.

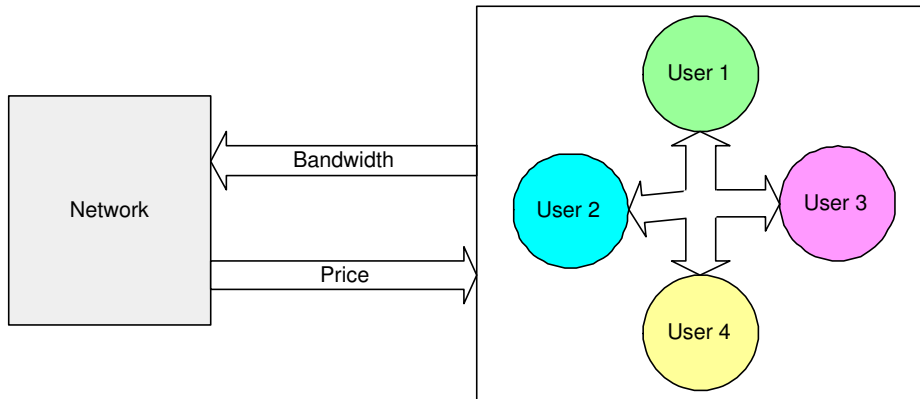


Figure 2.1: Network-User system interaction.

allocate the bandwidth among themselves by modifying their call request rates.

Fig. 2.1 illustrates this interaction between the network and users mediated through the exchange of price and bandwidth information. In reality, this exchange will pass through several iterations before attaining optimality. While the next chapter deals with rate optimization, we present price optimization in chapter 5.

2.1 Model Formulation

We consider a multi-link, multi-service connection oriented network composed of regular or virtual circuits. The links are labeled $1, 2, \dots, J$, with link j having a capacity of C_j circuits. We denote the set of all routes by R where a route r corresponds to a vector $a_{.,r}$. An element $a_{j,r}$ of this vector indicates the number of circuits required by a call through route r from link j ². The routing matrix is assumed to be fixed and is defined as $A = (A_{rj}, r \in R, j \in J)$ where $A_{rj} = 1$ if user r employs link j and $A_{rj} = 0$ otherwise.

Calls requesting route r arrive as a Poisson process of rate λ_r . A call requires one circuit in each of the links present in its route r implying $a_{j,r} = 1 \forall j \in r$. The call is connected and held for the call duration if there is a circuit free in all the links on its

²Corresponding to our abstract notion of “users”, we also use the term “calls” generically, to signify connections, transactions, bursts or flows generated by one of our users, and that occupy a virtual circuit for a period of time.

route, else it is blocked and lost. The call holding times are independent and identically distributed random variables with a mean duration of μ_r .

The calls from a user r pass through route r and request a QoS of θ_r . The network charges user r a rate of $\$M_r$ per circuit per unit-time. Since the market price is user specific, this is an example of third-degree price discrimination. The utility function $U_r(\theta_r)$ of user r is assumed to be continuous and strictly concave and depends only on its QoS θ_r . In practice, each user will try to choose its performance metric θ so as to maximize its net benefit (i.e., utility minus cost), $U_r(\theta_r) - M_r\theta_r$. The optimal QoS is thus obtained as

$$\theta_r^*(M_r) = \arg \max_{\theta_r} U_r(\theta_r) - M_r\theta_r = (U')^{-1}(M_r). \quad (2.1)$$

Depending on traffic conditions, individual calls from a user r would issue a connection set-up request at a rate λ_r . The QoS function $\mathcal{F}(\cdot)$ expresses the dependence of the performance metric (e.g., average number of circuits held, blocking probability, overflow probability, etc.) on the network capacity and the connection request rate vector $\Lambda \equiv [\lambda_1, \dots, \lambda_{|R|}]^T$ where $|R|$ denotes the cardinality of the route set R . For notational simplicity, the cardinality of R is often denoted by N in this dissertation. Information about the call arrival rate vector Λ is considered common knowledge and may be disseminated either by the users or the network.

Now suppose that the current arrival rate vector is $\Lambda(k)$. The r^{th} user will choose a new arrival rate $\lambda_r(k+1)$ so as to attain its desired θ_r^* using the current arrival rates of the other users³, $\lambda_j(k)$ for all $j \neq r$. The r^{th} user will choose $\lambda_r(k+1)$ that satisfies

$$\theta_r^* = \mathcal{F}_r(\lambda_r(k+1), \Lambda_{-r}(k), \mathcal{C}),$$

where $(\lambda_r(k+1), \Lambda_{-r}(k))$ represents a vector equal to $\Lambda(k)$ except that the r^{th} entry is $\lambda_r(k+1)$ instead of $\lambda_r(k)$. If no such $\lambda_r(k+1)$ exists in $S_r \equiv [0, \lambda_r^{\text{max}}]$, the user will instead choose the maximum possible rate $\lambda_r(k+1) = \lambda_r^{\text{max}}$. Therefore, a more compact expression is

$$\lambda_r(k+1) = \min\{(\tilde{\mathcal{F}}_r)^{-1}(\theta_r^*), \lambda_r^{\text{max}}\} \quad \forall r \in R \quad (2.2)$$

where $\tilde{\mathcal{F}}_r(\cdot) \equiv \mathcal{F}_r((\cdot, \Lambda_{-r}(k)), \mathcal{C})$. We define the set $S \equiv \times_{r \in R} S_r$, which is a nonempty, compact, convex subset of $\mathfrak{R}^{|R|}$. When the function $\tilde{\mathcal{F}} \equiv [\tilde{\mathcal{F}}_1, \dots, \tilde{\mathcal{F}}_{|R|}]^T$ is a continuous

³This assumption may not always hold, as in the context of imperfect information or arbitrary delays.

mapping from S to S , Brouwer's fixed point theorem [4] establishes the existence of a fixed point $\Lambda^{eqm} \in S$,

$$\lambda_r^{eqm} = \min\{(\tilde{\mathcal{F}}_r)^{-1}(\theta_r^*), \lambda_r^{max}\} \quad \forall r \in R. \quad (2.3)$$

In general, these fixed points may not be unique. The users thus allocate the available bandwidth among themselves in a distributed fashion through a noncooperative game. This bandwidth allocation game (*BAG*) can be formally expressed as

$$(BAG) \quad \max_{\lambda_r \in S_r} U_r(\theta_r(\lambda_r, \Lambda_{-r})) - M_r(\theta_r(\lambda_r, \Lambda_{-r})) \quad (2.4)$$

Furthermore, the convexity and compactness of the decision set S_r coupled with the concavity of the utility functions $U_r(\cdot)$ establish that the above equilibrium (2.3) is also a Nash Equilibrium Point (NEP) for the bandwidth allocation game. At the Nash equilibrium, no user can improve her QoS by unilaterally varying her call arrival rates.

The rate information of (2.2) may be exchanged with other users for a perfect information game as in the next chapter. Alternatively, it could be inferred by other users from their measured blocking probability as in chapter 4. We envisage these interactions to be performed by intelligent agents [28] to hide the complexity of difficult tasks from humans as well as provide rapid response to network variations. When all the agents employ some equation based model of the blocking probability, it may be possible to reduce the impact of user adaptation on network performance. Users will then exchange their rate vectors and carry out the computations without modifying their real arrival rates. Once the rate vectors stabilize to their equilibrium value (2.3), the users could then change their call request rates to attain the NEP. Undesirable performance fluctuations due to uncoordinated user adaptations can thus be avoided.

Chapter 3

Bandwidth Sharing in Circuit Switched Networks

In this chapter, we look at a specific application of our general model namely bandwidth sharing in circuit-switched networks fed by Poisson traffic. Unlike packet switched IP networks, their end-to-end behavior is well understood. Explicit formulae exist for blocking probabilities experienced by a single link. In the case of multi-link networks, techniques like the Erlang fixed approximation provide computationally feasible methods to estimate blocking.

Consider a mutli-link circuit switched network with $1, \dots, J$ links with capacity \mathcal{C} circuits shared by a pool of users. The users could be classes having the different origin - destination pair Label Switched Paths (LSPs) in a MPLS network or processes sharing pathways in a microprocessor. The QoS parameter Θ of concern here is the average number of circuits captured by each user. For the r^{th} user with a blocking probability of $B_r(\Lambda)$, the net arrival rate per user is

$$\lambda_r(1 - B_r(\Lambda)).$$

This blocking probability may be obtained analytically or estimated through online measurements. By Little's formula, the mean number of occupied circuits for the r^{th} user

is

$$\theta_r(\Lambda) \equiv \frac{\lambda_r}{\mu_r}(1 - B_r(\Lambda))$$

where we have explicitly shown the dependence of θ_r on all of the arrival rates Λ . The arrival rates reach the Nash equilibrium (2.3)

$$\lambda_r^{eqm} = \min\left\{\frac{\mu_r\theta_r^*}{1 - B_r(\Lambda^{eqm})}, \lambda_r^{max}\right\} \quad \forall r \in R. \quad (3.1)$$

The equilibrium allocation Θ^{eqm} can be expressed as

$$\begin{aligned} \theta_r^{eqm} &= \frac{\lambda_r^{eqm}}{\mu_r}(1 - B_r(\Lambda^{eqm})) \\ &= \min\left\{\frac{\mu_r\theta_r^*}{1 - B_r(\Lambda^{eqm})} \cdot \frac{1 - B_r(\Lambda^{eqm})}{\mu_r}, \lambda_r^{max} \cdot \frac{1 - B_r(\Lambda^{eqm})}{\mu_r}\right\} \\ &= \min\left\{\theta_r^*, \frac{\lambda_r^{max}(1 - B_r(\Lambda^{eqm}))}{\mu_r}\right\} \\ &\leq \theta_r^* \end{aligned} \quad (3.2)$$

since $0 \leq B_r(\cdot) \leq 1, \mu_r \geq 0$. The final allocation is suboptimal compared to the desired allocation, $\theta_r^{eqm} < \theta_r^*$ when

$$\lambda_r^{max} < \frac{\theta_r^* \mu_r}{1 - B_r(\Lambda^{eqm})}. \quad (3.3)$$

This occurs whenever the rate vector is stuck at the boundary due to the constraint box being too small or the capacity being insufficient to meet the demand. Thus for all the users, equilibrium allocation is never greater than the desired allocation. This fairness criterion ensures that no user receives less than he desired because of somebody receiving more than his due.

In the rest of this chapter, we extend the model from a single link to any general multi-link network. The implications in using upper bounds for approximating blocking probabilities is also considered. The real world applications include circuit switched optical networks, LSPs in MPLS and ATM networks. Virtual Private Networks built over the existing public Internet infrastructure are another potential candidate.

3.1 Erlang System

We begin by looking into the case of a single pool of resources, such as a single link, single virtual path or single class of traffic. The system of users and the network is

modeled as a stationary multi-class M/GI/K/K queue with total traffic intensity

$$\rho = \sum_{n=1}^N \frac{\lambda_n}{\mu_n}.$$

An interesting application of this model would be in the case of optical networks switching wavelengths using MPLS [33]. The aggregate and per-class connection blocking probability in steady state is then given by Erlang's formula [48],

$$\mathcal{E}(\rho, K) \equiv \frac{\rho^K / K!}{\sum_{k=0}^K \rho^k / k!}. \quad (3.4)$$

We obtain the classical Erlang system when the user arrival rates follow a Poisson process. The system of users and the network can then be modeled as a stationary multirate M/GI/C/C queue with total traffic intensity

$$\rho = \sum_{r \in R} \frac{\lambda_r}{\mu_r}.$$

The aggregate and per-class connection blocking probability in steady state is given by Erlang's formula [48],

$$\mathcal{E}(\rho, C) \equiv \frac{\rho^C / C!}{\sum_{r=0}^C \rho^r / r!}.$$

The user equilibrium call arrival rates (3.1) can then be computed as

$$\lambda_r^{eqm} = \min \left\{ \frac{\mu_r \theta_r^*}{1 - \mathcal{E}(\rho(\Lambda^{eqm}), C)}, \lambda_r^{max} \right\}. \quad (3.5)$$

In general, the equilibrium rate vector Λ^{eqm} would be dependent on the starting point Λ^0 . The presence of multiple Nash equilibria would make it difficult to predict the equilibrium rate and bandwidth allocation vectors. We now present a result to ensure the uniqueness of the equilibrium when there is no restriction on the user call request rates.

Lemma 1.

$$\rho(1 - \mathcal{E}(\rho, C)) \leq C \quad (3.6)$$

Proof. From [16, 29], the Erlang blocking probability \mathcal{E} is convex for all real positive values of C . Expressing $\mathcal{E}(\rho, C)$ simply as \mathcal{E} in order to simplify the notation,

$$\frac{d^2 \mathcal{E}}{d\rho^2} \geq 0$$

The slope of the Erlang blocking probability is thus monotonically nondecreasing.

$$\begin{aligned}
\frac{d\mathcal{E}}{d\rho} &= \left(\frac{C}{\rho} - 1\right)\mathcal{E} + \mathcal{E}^2 \\
&= \frac{C - \rho(1 - \mathcal{E})}{\rho} \mathcal{E} \\
&= \frac{C + \sum_{j=1}^{C-1} \left(\frac{C}{j!} - \frac{1}{(j-1)!}\right) \rho^j}{\rho \sum_{j=0}^C \frac{\rho^j}{j!}} \frac{\frac{\rho^C}{C!}}{\sum_{j=0}^C \frac{\rho^j}{j!}} \\
&= \frac{C + \sum_{j=1}^{C-1} \left(\frac{C}{j!} - \frac{1}{(j-1)!}\right) \rho^j}{\sum_{j=0}^C \frac{\rho^j}{j!}} \frac{\frac{\rho^{C-1}}{C!}}{\sum_{j=0}^C \frac{\rho^j}{j!}} \\
\implies \frac{d\mathcal{E}(0, C)}{d\rho} &= 0
\end{aligned}$$

From the monotonicity property,

$$\begin{aligned}
\frac{d\mathcal{E}}{d\rho} &\geq 0 \quad \forall \quad \rho \in [0, \infty). \\
\left(\frac{C}{\rho} - 1\right)\mathcal{E} + \mathcal{E}^2 &\geq 0 \\
\left(\frac{C}{\rho} - 1\right) &\geq -\mathcal{E} \\
1 - \mathcal{E} &\leq \frac{C}{\rho} \\
\rho(1 - \mathcal{E}) &\leq C
\end{aligned}$$

□

When there are no restrictions on arrival rates ($\lambda_r^{max} = \infty \quad \forall r \in R$), the equilibrium (3.5) becomes,

$$\lambda_r^{eqm} = \frac{\mu_r \theta_r^*}{1 - \mathcal{E}(\rho^{eqm}, C)} \quad \forall r \in R \tag{3.7}$$

$$\sum_{r \in R} \frac{\lambda_r^{eqm}}{\mu_r} = \frac{\sum_{r \in R} \theta_r^*}{1 - \mathcal{E}(\rho^{eqm}, C)}$$

$$\rho^{eqm}(1 - \mathcal{E}(\rho^{eqm}, C)) = \sum_{r \in R} \theta_r^*. \quad (3.8)$$

As a consequence, *Lemma 1* indicates that (3.8) has a solution only if the total demand is no greater than the total supply:

$$\sum_{r=1}^N \theta_r^* \leq C \quad (3.9)$$

Theorem 1. *Under the stability assumption (3.9) and no restrictions on the maximal arrival rates, there exists a unique nonnegative Nash equilibrium for the noncooperative bandwidth allocation game (2.4).*

Proof. From (3.7) we obtain the following equivalent system of equations by expressing λ_r in terms of λ_1 as

$$\lambda_r = \frac{\mu_r \theta_r}{\mu_1 \theta_1} \lambda_1 \quad \forall r \in R \setminus \{1\}. \quad (3.10)$$

Summation leads to

$$\rho(1 - \mathcal{E}(\rho)) = \sum_{r=1}^N \theta_r. \quad (3.11)$$

Λ^{eqm} can be reconstructed from ρ^{eqm} as

$$\begin{aligned} \rho^{eqm} &= \sum_{r=1}^N \frac{\lambda_r^{eqm}}{\mu_r} \\ &= \frac{\lambda_1^{eqm}}{\mu_1 \theta_1^*} \sum_{r=1}^N \theta_r^* \\ \implies \lambda_1^{eqm} &= \frac{\mu_1 \theta_1^* \rho^{eqm}}{\sum_{r=1}^N \theta_r^*}. \end{aligned}$$

Once λ_1^{eqm} is computed, the remaining call request rates can be computed from (3.10) .

Thus solving the $|R|$ dimensional system has been reduced to finding the roots of a one

dimensional problem. Expanding the terms of $\mathcal{E}(\rho)$, (3.11) can be rewritten as a C^{th} degree polynomial with real coefficients:

$$\sum_{r=0}^C a_r \rho^r = 0$$

where

$$a_r \equiv \frac{1}{r!} (D - r) \quad r = 0, \dots, C$$

and

$$D \equiv \sum_{r \in R} \theta_r^*, \quad 0 \leq D \leq C$$

Let $j = \lfloor D \rfloor$ and note that $0 \leq j \leq C$. Then the coefficients a_0 to a_j are positive and a_{j+1} to a_C are negative. There is only a single sign change of coefficients at $j + 1$. For a polynomial with real coefficients, Descartes' Sign Rule states that the number of positive real roots is equal to the number of variations in the sign of the coefficients or is less than that by an even integer. Thus, there exists a unique positive root to (3.11) and, by the nonnegativity of μ_r, θ_r , a unique nonnegative Nash equilibrium. \square

Fig. 3.5 illustrates how the roots of the above fixed point vary as the user demand is changed. As proved in the lemma, the equation has a solution only when the user demand can be satisfied by the network.

3.2 Speeding up Computations

The Erlang-B formula computes the blocking probability through a recursive procedure. The total number of recursion calls is equal to the number of circuits present in the link. This could be prohibitive for large capacities thereby hampering the deployment




Figure 3.1: Roots for Erlang B system.

of our rate control algorithm in real time scenarios. Instead, a *nonrecursive formula* is preferable, which could approximate and provide an upper bound for the blocking probability. While the performance of such a bound is independent of the capacity C , thus leading to fewer computations, their nonrecursive nature does away with the storage of previously computed values. We study two such bounds here, one proposed by Farago [11] and another by Martinez [6] which, henceforth in this paper, will be referred to as the Farago bound and Martinez bound respectively.

First, we present a result below on the effect of using upper bounds on the equilibrium arrival rates.

Lemma 2. *Consider the fixed point systems*

$$x = f(x) \tag{3.12}$$

$$x = g(x) \tag{3.13}$$

where $f(\cdot)$ and $g(\cdot)$ are continuous and nondecreasing functions of scalar x . Furthermore,




Figure 3.2: Graphical illustration of Lemma 2.

$g(\cdot)$ is an upper bound of $f(\cdot)$, $f(x) \leq g(x) \forall x \in \mathfrak{R}$. Let x^* and x^{**} be the fixed points of eqns. 3.12 and 3.13 respectively. Then x^{**} is an upper bound of x^* , $x^* \leq x^{**}$.

Proof. First $g(x^*) \geq f(x^*) = x^*$. Since $g(\cdot)$ is nondecreasing, it will therefore cross the line $y = x$ at a point $x^{**} \geq x^*$, see Fig. 3.2.

□

Thus, replacing blocking probabilities with their upper bounds leads to an equilibrium where the arrival rates are no lower than those with the true blocking probabilities. However, the equilibrium demands on the users are not affected when the network has adequate capacity and call request rates are unrestricted.

3.2.1 Farago Bound

We now detail the implications of using the Farago bound in our bandwidth game. Following Farago's convention, we define the time varying instantaneous bandwidth demand

of a traffic flow as ξ_t . The set of active flows is denoted by $A_t = \{\xi_t^{[r]} \mid \xi_t^{[r]} > 0\}$ with indices $r \in \langle A_t \rangle$ where $\langle A_t \rangle = \{r \mid \xi_t^{[r]} > 0\}$. The offered load $\phi(t)$ to the link is a function of all the active calls through it,

$$\phi(t) = \sum_{r \in \langle A_t \rangle} \xi_t^{[r]}.$$

Here we make the assumption that the offered load is the sum of the individual active flow bandwidth demands. Given the expected value of the offered load $F_t = E[\phi(A_t)]$ and the link capacity C satisfying the stability condition (3.9), the link blocking probability is bounded as

$$P(\phi(A_t) \geq C) \leq \left(\frac{F_t}{C}\right)^C e^{C-F_t}$$

Note that the bound tends to unity as the demand approaches the resource capacity. Further, the bound is meaningful only when the traffic intensity is less than C . The fixed point relation (3.1) thus reduces to

$$\lambda_r^{eqm} = \min \left\{ \frac{\mu_r \theta_r^*}{1 - B(\Lambda^{eqm}, C)}, \lambda_r^{max} \right\}, \quad (3.14)$$

where

$$B(\Lambda^{eqm}, C) = \left(\frac{\rho^{eqm}}{C}\right)^C e^{C-\rho^{eqm}}$$

Theorem 2. *Under the assumptions of stability and unrestricted maximal rates, the equilibrium arrival rates of the Farago system (3.14) is an upper bound on the equilibrium arrival rates obtained from (3.5).*

Proof. By definition, $B(\rho)$ is an upper bound of the Erlang blocking probability. From lemma 1, the Erlang-B blocking is a monotonic nondecreasing function. The slope of the Farago bound can be computed as

$$\begin{aligned} B'(\rho) &= \left(\frac{C}{\rho} - 1\right)B(\rho) \\ &\geq 0 \quad 0 \leq \rho \leq C, \end{aligned}$$

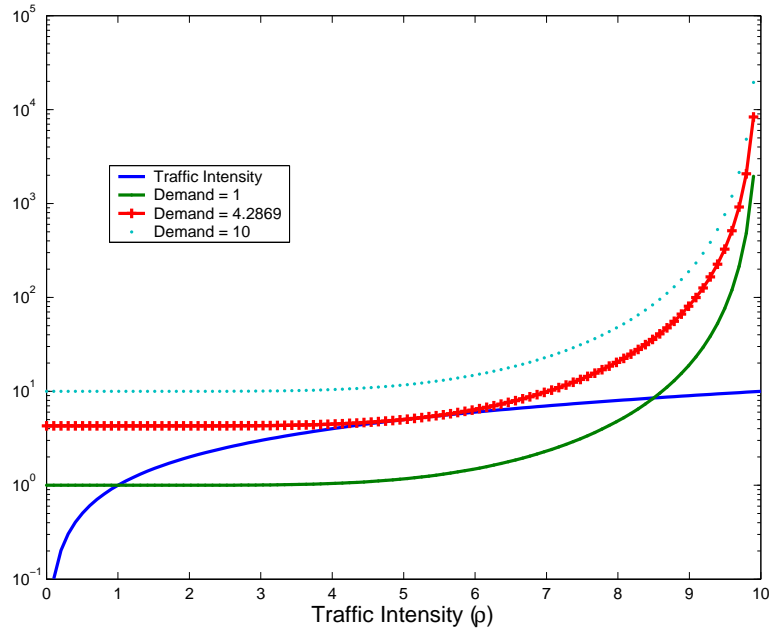


Figure 3.3: Roots of Farago system.

thereby satisfying the conditions of Lemma 2.

□

In the absence of rate constraints, there may exist none, one or two solutions to the fixed point (3.14). Fig. 3.3 illustrates a saddle point bifurcation for $C = 10$. The bifurcation occurs at the critical demand D_b which satisfies the condition that demand curve is tangent to the graph of the traffic intensity. The following conditions involving the zeroth and first derivative are satisfied, namely,

$$\begin{aligned} \frac{D_b}{1 - B(\rho_b)} &= \rho \\ \frac{d}{d\rho} \left(\frac{D_b}{1 - B(\rho_b)} \right) &= 1. \end{aligned}$$

Simplifying,

$$1 - B(\rho_b) = \frac{D_b}{\rho} \quad (3.15)$$

$$B'(\rho_b) = \frac{D_b}{\rho_b^2}. \quad (3.16)$$

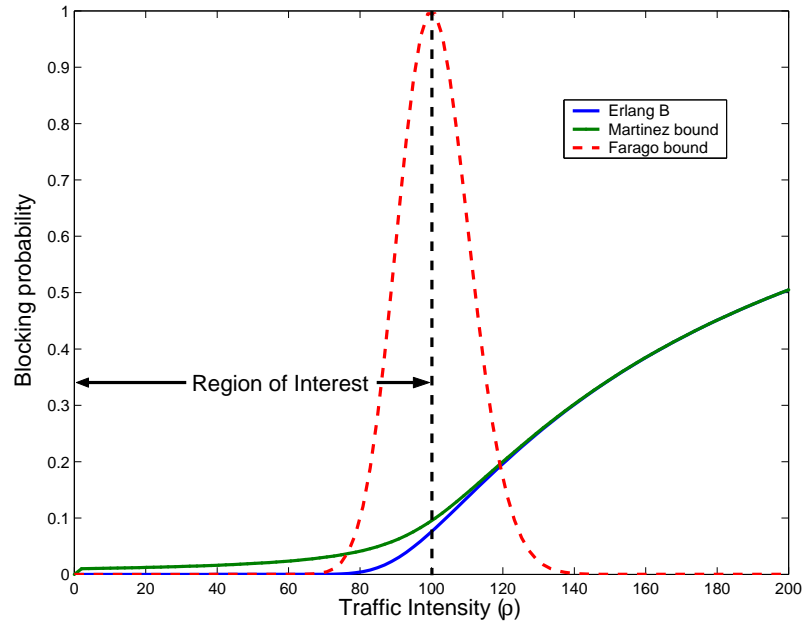


Figure 3.4: Comparison of bounds with blocking probability.

Dividing (3.16) by (3.15),

$$\frac{B'(\rho_b)}{1 - B(\rho_b)} = \frac{1}{\rho_b}$$

The critical traffic intensity (ρ_b) can then be obtained by solving

$$B(\rho_b)\rho_b - B(\rho_b)(\rho_b + 1) + 1 = 0 \quad (3.17)$$

The bifurcation demand D_b can then be computed from (3.15).

3.2.2 Martinez Bound

Martinez [6] introduced a strict upper bound on the blocking given by Erlang-B formula.

$$\mathcal{E}(\rho, C) < \frac{1}{2\rho}(\sqrt{(1 + C - \rho)^2 + 4\rho} - (1 + C - \rho))$$

A comparison of the Farago and Martinez bounds with respect to the Erlang blocking probability for $C = 100$ is illustrated in Fig. 3.4. While the Farago bound provides a better

approximation for low intensities, the Martinez bound tightly bounds the Erlang blocking probability for higher intensities. With the bound surrogating for the blocking, the fixed point (3.1) thus becomes

$$\lambda_r^{eqm} = \min \left\{ \frac{\mu_r \theta_r^*}{1 - B(\Lambda^{eqm}, C)}, \lambda_r^{max} \right\}, \quad (3.18)$$

where

$$B(\Lambda^{eqm}, C) = \frac{1}{2\rho^{eqm}} (\sqrt{(1 + C - \rho^{eqm})^2 + 4\rho^{eqm}} - (1 + C - \rho^{eqm}))$$

In the absence of rate constraints, there may exist one or no solutions to the fixed point (3.18).

Theorem 3. *In the absence of restrictions on maximal rates, the equilibrium arrival rates of the Martinez system (3.18) is an upper bound on the equilibrium arrival rates for the Erlang system.*

Proof. By definition, $B(\rho)$ is an upper bound of the Erlang blocking probability. From lemma 1, the Erlang-B blocking is a monotonic nondecreasing function. The slope of the Martinez bound can be computed as

$$B'(\rho) = \frac{-2\rho + (C + 1)(\sqrt{(1 + C - \rho)^2 + 4\rho} - (1 + C - \rho))}{2\rho^2 \sqrt{(1 + C - \rho)^2 + 4\rho}}$$

For nonnegative ρ, C , the denominator is always positive. We proceed to prove the nonnegativity of the numerator.

$$\begin{aligned} \frac{-C\rho}{(C + 1)^2} &\leq 0 \\ \Rightarrow 1 + \frac{\rho}{C + 1} \frac{-C}{C + 1} &\leq 1 \\ \Rightarrow 1 + \frac{\rho}{C + 1} \left(\frac{1}{C + 1} - 1 \right) &\leq 1 \\ \Rightarrow 1 + \frac{\rho}{(C + 1)^2} - \frac{\rho}{C + 1} &\leq 1 \\ \Rightarrow \frac{\rho}{(C + 1)^2} + \frac{C + 1 - \rho}{C + 1} &\leq 1 \end{aligned}$$

Multiplying by 4ρ and then adding $(1 + C - \rho)^2$,

$$\Rightarrow \left((1 + C - \rho) + \frac{2\rho}{C + 1} \right)^2 \leq (1 + C - \rho)^2 + 4\rho$$

Taking square roots,

$$\begin{aligned} &\Rightarrow \left((1 + C - \rho) + \frac{2\rho}{C + 1} \right) \leq \sqrt{(1 + C - \rho)^2 + 4\rho} \\ &\Rightarrow -2\rho + (C + 1)(\sqrt{(1 + C - \rho)^2 + 4\rho} - (1 + C - \rho)) \leq 0, \end{aligned}$$

thereby indicating the nonnegativity of the numerator. The Martinez bound is thus a monotonic nondecreasing function satisfying the conditions of lemma 2.

□

3.3 Product Form Networks

We now proceed to the general case of a multi-link network with fixed routing. Without loss of generality, we assume that the holding periods of calls on route r are identically distributed with unit mean. The routing matrix $A \equiv (a_{jr}, j = 1, \dots, J, r \in R)$ with a route r identified with a subset of the set of links $\{1, \dots, J\}$. Let $n_r(t)$ be the number of calls in progress at time t on route r , and define the vector $n(t) \equiv (n_r(t), r \in R)$. The continuous-time Markov chain $(n(t), t \geq 0)$ takes values in $\mathcal{S}(C) \equiv \{n \in \mathbb{Z}_+^R : An \leq C\}$ and has a unique stationary distribution given by

$$\pi(n) = G(C)^{-1} \prod_{r \in R} \frac{\lambda_r^{n_r}}{n_r!}, \quad n \in \mathcal{S}(C), \quad (3.19)$$

where $G(C)$ is the normalizing constant

$$G(C) = \sum_{n \in \mathcal{S}(C)} \prod_{r \in R} \frac{\lambda_r^{n_r}}{n_r!}. \quad (3.20)$$

This model has wide applicability in telephone networks, multiprocessor interconnection architectures, database structures, mobile radio and broadband packet networks. In modern computer communication and telephony networks, the circuits are often “virtual”

as in the case of a fixed proportion of the transmission capacity of a communication channel. Various generalizations may be incorporated into the above model to enlarge its scope. Note that if calls requesting route r arrive at rate λ_r/μ_r and have holding periods μ_r , then the resulting stationary distribution π remains unaltered from (3.19). Further if the arrival rate of calls in route r depends on the calls in progress n_r as in the Engset model with a finite source population, then the corresponding distribution is given by a minor variant of (3.19).

Note that the time average of the number of occupied circuits is provided by (3.19). This is in general not the same as the event average of the blocking observed by incoming calls. When the arrivals are obtained from a Poisson process, these two numbers coincide owing to the PASTA (Poisson Arrivals See Time Averages) property.

The stationary probability that a route r call is blocked is given by

$$L_r = 1 - \frac{G(C - Ae_r)}{G(C)},$$

where e_r is the unit vector from $S(C)$ describing just one call in progress on route r . Unfortunately, the computation of the normalizing constant $G(C)$ cannot usually be computed in polynomial time. This is because the number of routes $|R|$ may grow exponentially with the number of links J . Consider the trivial case where route r corresponds to a single link r with $C_j = C \ \forall \ j$. Then the number of routes equals the number of links ($|R| = J$, $A = I$). Clearly, the size of the state space $|S(C)| = \prod_{j=1}^J C_j = C^J$ illustrating the exponential growth. Thus it is imperative to develop approximations to reduce the computational complexity as well as provide deeper insights to the problem.

We now detail the celebrated *Erlang Fixed Point Approximation* for tackling this conundrum. For a loss network with fixed routing and a 0–1 routing matrix, let B_1, \dots, B_J be a solution of the system of equations

$$B_j = \mathcal{E}(\rho_j, C_j), \quad j = 1, \dots, J \quad (3.21)$$

$$\rho_j = \sum_r \lambda_r \Pi_{i \in r - \{j\}} (1 - B_i), \quad j = 1, \dots, J \quad (3.22)$$

where the function $\mathcal{E}(\cdot)$ is the Erlang formula (3.4). Kelly [20] showed that there exists a unique vector $(B_1, \dots, B_J) \in [0, 1]^J$ satisfying (3.21)-(3.22), which we term the *Erlang fixed point*. Then an approximation for the loss probability on route r is given by

$$L_r \cong 1 - \prod_i (1 - B_i)^{a_{ir}}, \quad r \in R \quad (3.23)$$

Employing the Erlang fixed point approximation, the Nash equilibrium of our multi-link, multi-user game is obtained as

$$\lambda_r^{eqm} = \min \left\{ \frac{\theta_r^*}{1 - L_r(\Lambda^{eqm}, C)}, \lambda_r^{max} \right\}. \quad (3.24)$$

The rationale behind the approximation is as follows. If the links in route r were blocked independently (which they clearly are not) and if B_j were the link-blocking probability, then L_r would be the route r blocking probability

$$L_r = 1 - \prod_{i \in r} (1 - B_i) = 1 - \prod_i (1 - B_i)^{a_{ir}}.$$

The traffic offered to link j would then be Poisson at rate ρ_j and the carried traffic on link j would be

$$\sum_r a_{jr} \lambda_r (1 - L_r) = \rho_j (1 - B_j).$$

The Erlang Fixed Point Approximation requires that the blocking probabilities (B_1, \dots, B_J) should be consistent with this level of carried traffic. The expression (3.22) is often referred to as the *reduced load* on link j .

Such an approximation was shown to be exact under the limiting regimes of moderate loading [20] and diverse routing [15]. The rates regulated by the Erlang fixed point approximation have the desired property that the corresponding bandwidth allocations remain feasible.

Theorem 4. *The bandwidth allocation vector Θ^{eqm} corresponding to the Nash equilibrium point under the Erlang fixed point approximation is feasible, i.e., $A\Theta^{eqm} < C$.*

Proof. Let the Nash equilibrium point of the user game be denoted as Λ^{eqm} . Define R_j as the set of all routes passing through link j . For a certain link j , the total bandwidth

allocated to all the routes would be

$$\begin{aligned}
\sum_{r \in R_j} \theta_r^{eqm} &= \sum_{r \in R_j} \lambda_r^* (1 - L_r) \\
&= \sum_{r \in R_j} \lambda_r^* \prod_{i \in r} (1 - B_i) \\
&= (1 - B_j) \sum_{r \in R_j} \lambda_r^* \prod_{i \in r - \{j\}} (1 - B_i) \\
&= (1 - B_j) \rho_j \\
&= (1 - \mathcal{E}(\rho_j, C_j)) \rho_j \\
&\leq C_j \quad \forall j = 1, \dots, J,
\end{aligned}$$

by lemma 1. □

3.4 Results and Discussion

We present the results of our investigation into the dynamics of a two-class game for resource sharing in a single link. Each user has a utility function of the form

$$U_r(\theta_r) = b_r \log(1 + \theta_r), \quad b_r > 0,$$

where we denote b_r as the utility coefficient of user r . Thus the optimal bandwidth of each user is given by

$$\theta_r^* = \frac{b_r}{M_r} - 1.$$

We can observe that the optimal bandwidth is thus proportional to the utility coefficient. Unless otherwise specified, the users analyzed were identical with $b = 10$, $\lambda^{max} = [20 \ 20]^T$ and $\mu = [1 \ 1]^T$.

While computing the arrival rates using the Erlang formula (3.5), we experimented with two termination criteria - one based on the deviation between successive call request rates ($\Delta\Lambda$) and another between successive bandwidth allocations ($\Delta\Theta$). Fig. 3.5 compares the termination criteria with respect to the total number of iterations required to attain equilibrium. While the number of iterations for the $\Delta\Lambda$ increase linearly with respect to a

decrease in the tolerance δ , the convergence under $\Delta\Theta$ criterion remains almost unchanged. This is due to the lower sensitivity of the bandwidth (θ) to changes in call request rates. All our calculations were thus terminated using the $\Delta\Theta$ benchmark.




Figure 3.5: Comparison of termination criteria.

3.4.1 Convergence of Arrival Rates

The convergence of the call arrival rates to the Nash equilibrium is influenced by the demand-supply gap and the size of the bounding box of maximal arrival rates. Fig. 3.6 illustrates a system where the total user demand $D = 9$ is less than the supply $C = 10$. The unrestricted Nash equilibrium $\Lambda^{eqm} = [11.01 \ 5.51]^T$ is infeasible when the maximal rates are too restrictive. Thus the arrival rates become stuck at the boundary of the box. However when the box size is increased, the system converges to its unrestricted equilibrium Λ^{eqm} . Increasing the maximum rates further has no effect on the convergence. There exists no unrestricted Nash equilibrium when the user demand $D = 12$ is more than the network capacity $C = 10$ as in Fig. 3.7. Thus the arrival rates eventually converge to the box boundary irrespective of the maximum arrival rates.




Figure 3.6: Arrival rates under resource surplus in a single link




Figure 3.7: Arrival rates under resource deficit in a single link

Uniqueness of Nash Equilibrium




Figure 3.8: Convergence to NEP.

For the single link scenario, the convergence of the user game to a unique equilibrium under sufficient capacity and maximal rates was proved in theorem 1. Fig. 3.8 shows the convergence of the Erlang fixed point iteration (3.5) for various starting values of the user arrival rates, Λ . The Λ state space was also exhaustively scanned for other prospective candidates. It was observed that all the iterations converged to the same Nash equilibrium point irrespective of their starting values.

For a sequence $\{x^r\}$ converging to x^* in \mathfrak{R}^n , the Q-rate convergence is linear if there exists $q > 0$ and $\beta \in (0, 1)$ such that for all r

$$\|x^r - x^*\| \leq q\beta^r.$$

Figure 3.9 indicates that the rate of convergence for (3.5) is linear. The linear rate of convergence and the absence of alternative equilibria indicate that the starting vector was chosen within the region of contraction.




Figure 3.9: Speed of convergence to NEP.

3.4.2 Effect of Bounds on Nash Equilibrium

Although employing upper bounds may lead to computational efficiency, Fig. 3.10 illustrates that the resulting equilibrium call request rates are higher than those obtained using the Erlang-B formula. Whenever the maximum arrival rates are low enough, these run the risk of being rendered infeasible. The iterations then are caught on the maximal rate boundary leading to a suboptimal bandwidth allocation (3.3).

3.4.3 Triangle Network

As an example of product form networks, we consider the three link, triangle network of Fig. 3.11. It possesses two user flows flowing through routes 1 and 2 respectively. Route 1 consists of links 1 and 3 while route 2 is composed of links 2 and 3. Thus $a_{11} = 1$, $a_{31} = 1$, $a_{22} = 1$, $a_{32} = 1$. The link capacity vector is $C = [10, 10, 20]^T$. Letting the service




Figure 3.10: Convergence to equilibrium rates.




Figure 3.11: Triangle network.

rate to be identically unity, the Erlang fixed point equations (3.21)-(3.22) become

$$\begin{aligned}
B_1 &= \mathcal{E}(\rho_1, C_1) \\
B_2 &= \mathcal{E}(\rho_2, C_2) \\
B_3 &= \mathcal{E}(\rho_3, C_3) \\
\rho_1 &= \lambda_1(1 - B_3) \\
\rho_2 &= \lambda_2(1 - B_3) \\
\rho_3 &= \lambda_1(1 - B_1) + \lambda_2(1 - B_2).
\end{aligned}$$

The initial reduced load vector $\rho(0)$ is chosen as $[1, 1, 1]$. After computing the new reduced load vector, we calculate the new B iterate. Once the iterations converge to obtain the final $[B_1^*, B_2^*, B_3^*]$, the route blocking probabilities are computed as

$$\begin{aligned}
L_1 &= 1 - (1 - B_1)(1 - B_3) \\
L_2 &= 1 - (1 - B_2)(1 - B_3)
\end{aligned}$$

where L_1, L_2 are the approximate blocking probabilities for routes 1 and 2 respectively.

As in the single link case, we are interested in the partitioning of the link capacities between competing users under various maximal rates and user demand characteristics. The bandwidth price is fixed at $M = [1 \ 1]^T$ and the iterations were initiated at $\Lambda(0) = [1 \ 1]^T$. First, we consider the resource surplus case wherein user utility coefficients are $b = [5 \ 10]^T$ leading to user demands $\Theta^* = [4 \ 9]^T$. A bandwidth surplus of 6, 1 and 7 occurs at links 1, 2 and 3 respectively. For a triangle network, the convergence of the user rates to the Nash equilibrium is illustrated in Fig. 3.12. The corresponding equilibrium bandwidth allocations Θ^{eqm} are shown in Fig. 3.13. It can be seen that the convergence of the user game is dependent on the maximal rates. When they are too restrictive as in the case of $\Lambda^{max} = [10 \ 10]^T$, the iterations are stuck at the boundary. This causes the allocation of user 2 to be below her desired value $\theta_2^* = 9$ even though the network has surplus capacity.

The more interesting case occurs when the user demands for bandwidth are more than that could be satisfied by the network. In chapter 6, we discuss how the network would modify its network price whenever user demand outstrips capacity. The user utility coefficients are fixed at $b = [10 \ 20]^T$ leading to user demands of $\Theta^* = [9 \ 19]^T$. User 2's appetite for bandwidth leads to a resource deficit of 9, 8 in links 2 and 3 respectively. This pits user 1 in direct competition to user 2 for the scarce resources at link 3. Since the




Figure 3.12: Arrival rates under resource surplus in a triangle network.

situation is clearly unsustainable, the user game does not converge to a finite equilibrium. The rates thus continue to increase till they hit the maximal rate boundary as indicated in Fig. 3.14. Also user 2 enhances her rates much more than user 1 due to her higher demand. From Fig. 3.15, it can be observed that the network tries to satisfy user 2 while being fair to user 1. Thus user 1 receives her desired 9 units of bandwidth with the rest allocated to user 2. As the maximal rates are increased, the network allocates more and more of the leftover capacity to user 2.




Figure 3.13: Bandwidth allocation under resource surplus in a triangle network.




Figure 3.14: Arrival rates under resource deficit in a triangle network.




Figure 3.15: Bandwidth allocation under resource deficit in a triangle network.

Chapter 4

Imperfect Information Regime

The existence of Nash equilibria in an environment populated by self-interested users rests on the assumption of perfect information. This requires every user to be aware of the state of all others at any point in time. Errors in measurement coupled with delay in information propagation can violate this condition. Users might also try to conceal their behavior so as to deny competitors of any unfair advantage. Another subtle cause for imperfection is the assumption made while deriving the Erlang blocking formula and the Erlang fixed point approximation. These require that the user rates be independent of one another. While this assumption will be satisfied at the Nash equilibrium Λ^{eqm} , that may not be the case when the users are far from the fixed point. User rates based on the observed blocking in the network are implicitly dependent on the behavior of others and hence stochastically not independent of each others' arrival rates. Hence any application of game theory based flow control strategies in a real world setting would have to address the issue of imperfect information. Alpcan et al. in [2] acknowledge the inherent restrictions of implementing cost functions in Internet-style networks and propose a scheme based on the variations in queuing delay of the individual user. In contrast, we derive our inspiration from optimal estimation based adaptive control techniques for stochastic dynamical systems [44, 45]. An analogous approach involving the Kalman filter was employed by Alouf et al. in [1] for the on-line estimation of dynamic multicast groups.

In this chapter, we extend our bandwidth allocation model developed for connection oriented networks in [31, 33] to encompass an imperfect information regime, character-

istic of real world settings. In [3], we propose a distributed adaptive algorithm to be used by each user to attain her desired optimum in the presence of uncertainty. The instantaneous bandwidth utilization observed while entering the network is employed to infer the system state through recursive estimations. Users then modify their arrival rates to maximize their individual utilities. We also develop two variants of the original algorithm to account for the curvilinearity of the input-output response. These adaptive control schemes were simulated for making comparisons of scalability and performance under various system parameters. Results indicate the algorithms converge even in the presence of uncertainty about the number of other players and their strategies. Applications of this algorithm include teletraffic, wireless and optical networks, enabling users to partition bandwidth without the need of a centralized synchronizing entity.

4.1 Erlang System

We consider a loss network scenario where N users compete for a finite amount of bandwidth K as expostulated in [31, 33]. Each user has a utility function $U_i(\theta)$ which is maximized at an optimal bandwidth θ_i^* . Requests for bandwidth arrive as a Poisson process with rate λ_i . Queuing is not a concern here as unfulfilled user requests depart the system immediately. The behavior of competing users can then be modelled as a noncooperative game wherein each player strives to attain her optimal θ_i^* by regulating her sending rate λ_i . Under the assumption that total demand for bandwidth is not greater than supply $\sum_{i=1}^N \theta_i^* \leq K$, the mean bandwidth allocated to the i^{th} user is given by Little's formula as

$$\theta_i(\Lambda) = \frac{\lambda_i}{\mu_i} (1 - \mathcal{E}(\rho(\Lambda))). \quad (4.1)$$

Here Λ is the user arrival rate vector $[\lambda_1, \dots, \lambda_N]$ and μ_i is the exponential service rate for user i . The traffic intensity is defined as $\rho \equiv \sum_{i=1}^N \frac{\lambda_i}{\mu_i}$ and blocking probability as $\mathcal{E}(\rho, K) \equiv \frac{\rho^K / K!}{\sum_{k=0}^K \rho^k / k!}$. As derived in [31], in a perfect information regime where each user is aware of every other's sending rates, the user arrival rate vector converges to the non-cooperative Nash equilibrium Λ^{eqm}

$$\lambda_i^* \equiv \frac{\mu_i \theta_i^*}{1 - \mathcal{E}(\rho(\Lambda^{eqm}))} \quad \forall \quad i = 1, \dots, N \quad (4.2)$$

The postulation of perfect information is often violated in real networks. The requirement that all the players need to be updated with the current information simultaneously places

an enormous onus on the system. A distributed version of such an update protocol would involve $O(N^2)$ packets to be exchanged. Centralized broadcast of updates can reduce this to $O(N)$ but at the cost of increased synchronicity. Propagation and queuing delays cause staleness of information when user behaviors are highly variable. Internet connections are very often bursty and short lived thereby forcing us to examine the less tractable imperfect information scenario.

4.2 Feedback Based Rate Control

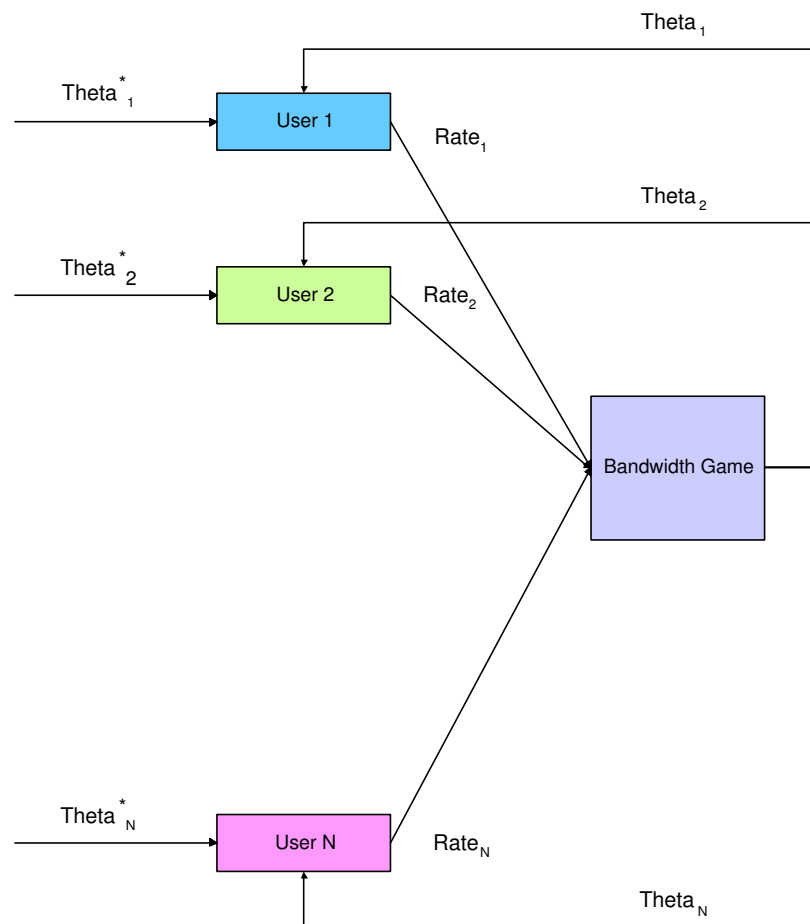


Figure 4.1: Feedback based user rate adaptation.

We now detail an observer based control scheme for the bandwidth allocation model system defined above. Under imperfect information, each user i is aware of only her individual tuple of variables - $(\mu_i, \lambda_i, \theta_i(k))$. While the service rate μ_i is characteristic to the user, the mean bandwidth consumed θ_i can be ascertained from its instantaneous value $\theta_i(k)$ at the k^{th} time step. Since the information about the available bandwidth is obtained only when a new arrival enters the network, the system is modeled using difference equations. Each user sees the network as a dynamical system evolving in time with ρ as the system state. She tries to estimate the state using observations $\theta_i(k)$ which she employs to compute an observer based feedback control λ_i designed to attain her optimal bandwidth θ_i^* . This feedback mechanism is illustrated in Fig. 4.1. The effect of other users on the system is modeled as noise which can be observed as perturbations of ρ from the hypothesized value. The state equation thus becomes

$$\rho(k+1) = \rho(k) - \frac{\lambda_i(k)}{\mu} + \frac{\lambda_i(k+1)}{\mu} + v(\rho(k)) \quad (4.3)$$

where $v(\rho(k))$ is the state dependent noise. Note that the noise is not statistically independent and hence cannot be modeled as Gaussian. In the absence of estimation errors due to uncertainty, the measured output can be related to the state and control using (4.1). The user calculates her new feedback control using the estimate of state $\hat{\rho}(k+1)$,

$$\lambda_i^{k+1} = \frac{\alpha_i(k+1)\theta_i^*\mu_i}{1 - \mathcal{E}(\hat{\rho}(k+1))} \quad (4.4)$$

Since α_i determines how fast each user tries to capture her optimal bandwidth, it can be considered as an indicator of user ‘‘aggressiveness’’. We thus define

$$\alpha_i(k) \equiv 1 - (\beta_i)^k, \quad 0 \leq \beta_i < 1. \quad (4.5)$$

In a perfect information scenario, the user would eventually attain her optimal bandwidth as $\lim_{i \rightarrow \infty} \alpha_i = 1$.

If the solution of the fixed point equation (4.2) is known to all the users, the user rates get decoupled. Thus the players move to the Nash equilibrium by the modified state equation

$$\lambda_i(k+1) = \lambda_i(k) + \eta(\lambda_i^* - \lambda_i(k)), \quad 0 < \eta \leq 1 \quad (4.6)$$

However in a decentralized game, each user is only aware of the dynamics of his individual state and observer equations. Thus the final equilibrium if attained may not be the Nash Equilibrium Point (NEP) as above.

4.3 Dynamic Estimation

The key ingredient in calculating an observer based feedback control is the estimation algorithm which provides an accurate estimate of the system state with minimal computational and storage requirements. The user collects observations $\{\theta(0), \theta(1), \dots, \theta(k)\}$ which she employs to estimate the time varying system state ρ^k . A recursive estimation procedure is vital in reducing the observation history to be maintained at each point in time. We make a first order approximation of the relationship between the inputs and measured outputs by assuming θ to be linear in ρ i.e $\theta(k) \simeq h(k)\rho(k) + w(k)$ where $h(k)$ is the design parameter and $w(k)$ is the unknown noise value. Each user then strives to reduce her least squares error

$$L(\rho, k) = \frac{1}{2k} \sum_{j=1}^k \gamma^{k-j} (\theta(j) - h(j)\rho(j))^2 \quad (4.7)$$

for the weight $0 \leq \gamma < 1$. Eqn. (4.7) exponentially de-weights past measurements indicating that greater importance is placed on current measurements. $L(\rho, k)$ is minimized by the classic Recursive Least Squares (RLS) algorithm which produces a time-varying estimate of the system state ρ [44] as

$$\hat{\rho}(k) = \hat{\rho}(k-1) - K_k(\lambda(k)\hat{\rho}(k-1) - \theta(k)) \quad (4.8)$$

where $K_k = P_k \lambda(k)$ and $P_k = \frac{P_{k-1}}{\gamma} (1 - \frac{P_{k-1}^2 (\lambda(k))^2}{\gamma + (\lambda(k))^2 P_{k-1}})$. The algorithm is initiated with $\hat{\rho}(0) = 0$, $P_0 \gg 0$.

4.4 Rate Adaptation Algorithms

We detail three algorithms below which would then be compared with respect to their convergence and accuracy among others.

4.4.1 Original Algorithm

This method estimates system state using (4.8) while the control is calculated by approximating (4.4) as

$$\lambda_i(k+1) = \frac{\alpha_i(k+1)\theta_i^* \mu_i}{1 - \mathcal{E}(\hat{\rho}(k))} \quad (4.9)$$

We denote it as the ‘‘original algorithm’’.




Figure 4.2: Comparison of measured and estimated outputs for 2 users.

4.4.2 Logarithmic Variant

In microeconomics, the relationship between an economic output (y) and its inputs (x) is often described by a Cobb-Douglas type of production function $y = C \prod_{i=1}^n x_i^{a_i}$. Thus there exists a linear relationship between the logarithmic values of input and output namely $\log y = \log C + \sum_{i=1}^n a_i \log x_i$. The correspondence between user inputs and observations for a two player scenario was analyzed as follows. For symmetric users, we collected multiple input-output data and computed the exponents for equation $\theta = C\lambda_1^{a_1}\lambda_2^{a_2}$ by least squares fitting.

Fig. 4.2 illustrates the sup-norm of the error between the measured and estimated outputs for various values of K . It is clear that the input-output relationship can be approximated by a log-linear one. The state estimation equation can then be rewritten as

$$\log \hat{\rho}(k) = \log \hat{\rho}(k-1) - K_k(\lambda(k) \log \hat{\rho}(k-1) - \log \theta(k))$$

The time varying feedback control is computed using (4.9). We denote the above algorithm as the “logarithmic variant”.

4.4.3 Newton-Raphson Variant

Another approximation of (4.4) could be carried out by ignoring the effect of the external noise $v(\rho(k))$ on the evolution of system state. Thus

$$\rho(k+1) = \hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\lambda(k+1)}{\mu} \quad (4.10)$$

Substituting (4.10) in (4.4), we obtain

$$\lambda(k+1) = \frac{\alpha(k+1)\theta^*\mu}{1 - \mathcal{E}(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\lambda(k+1)}{\mu})}$$

The task of computing the new control thus reduces to finding the root of $g(\lambda(k+1))$

$$g(\lambda(k+1)) \equiv \lambda(k+1) - \frac{\alpha(k+1)\theta^*\mu}{1 - \mathcal{E}(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\lambda(k+1)}{\mu})}$$

We employ the damped Newton-Raphson method for this purpose. Denoting $\delta(n)$ as the n^{th} iterate of $\lambda(k+1)$ and κ ($0 \leq \kappa < 1$) as the damping coefficient, the iterations are

$$\delta(n+1) = \delta(n) - \kappa g'(\delta(n))^{-1}g(\delta(n)) \quad (4.11)$$

where

$$g'(\delta(n)) = 1 - \alpha(k+1)\theta^* \left[\left(\frac{K}{\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\delta(n)}{\mu}} - 1 \right) \mathcal{E}(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\delta(n)}{\mu}) + \mathcal{E}^2(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\delta(n)}{\mu}) \right] / \left(1 - \mathcal{E}(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\delta(n)}{\mu}) \right)^2$$

and

$$g(\delta(n)) = \delta(n) - \frac{\alpha(k+1)\theta^*\mu}{1 - \mathcal{E}(\hat{\rho}(k) - \frac{\lambda(k)}{\mu} + \frac{\delta(n)}{\mu})}$$

Starting from $\delta(0) = \frac{\alpha(k+1)\theta^*\mu}{1 - \mathcal{E}(\hat{\rho}(k))}$, we compute the successive approximations of $\lambda(k+1)$ using (4.11). Termination of the iterations is contingent on the condition of the error falling below the tolerance threshold ξ i.e $|\delta(n+1) - \delta(n)| \leq \xi$. The ‘‘Newton-Raphson variant’’ thus estimates system state using (4.10) and computes the corresponding control by (4.11).

4.4.4 Simulation Results

Computer simulation was employed to investigate the behavior of the algorithm under several scenarios. For ease of analysis, the system considered was composed of a homogenous population of identical users. Unless mentioned otherwise, the default values used throughout the simulation were: $\theta^* = 20$, $\mu = 1$, $\beta = 0.1$, $P_0 = 10^3$, $\xi = 10^{-3}$, $\kappa = 0.1$. Each different scenario was executed for 1000 repetitions to obtain statistically significant results. For stochastic optimization methods like the ones used above, measurement noise makes it unrealistic to expect the algorithms to converge to a single value. We thus employed the following stopping criterion in our experiments. The algorithm is presumed to have stabilized at iteration n when 10 successive values of the system state are at most 0.1 apart from each other

$$|\rho(n) - \rho(n - j)| \leq 0.1 \quad \forall \quad 1 \leq j \leq 10$$

The maximum number of iterations until convergence was set to 10^5 beyond which it was terminated and considered as not converging.

The cardinal outcome of the experiments was the convergence of the algorithms in a noisy environment, ruling out the occurrence of system instabilities such as oscillations and finite time singularities. This is impressive considering the fact that each user pursues her utility optimization oblivious to the number of competitors or the algorithms adopted by them to attain their objectives. Further the equilibria under uncertainty are close to the Nash equilibrium in a perfect information regime. Other characteristics of the algorithms are detailed below.

4.4.5 Scalability

The scalability of an algorithm determines the speed at which it stabilizes the system for increasing number of users. The original algorithm does not scale well and does not converge when the number of users are greater than 5. Further as seen in Fig. 4.5, the number of iterations required for convergence are two orders of magnitude more than the two variants. This instability thus rules out its deployment in a real world setting forcing us to exclude it from further consideration. The logarithmic and Newton-Raphson variants scale well with N and are compared in Fig. 4.3 indicating the superiority of the former over the latter. The effect of other users on the constitution of system state becomes prominent

as the number of users proliferate. This affects the validity of assumption (4.10) and the scalability of the Newton-Raphson version.




Figure 4.3: Iterations vs Number of users: Logarithmic and Newton-Raphson.

4.4.6 Accuracy

The effect of noise and delay in an imperfect regime is to move the system away from the non-cooperative Nash equilibrium. Hence a suitable metric to quantify the superiority of the algorithms would be the deviation of their equilibria (Λ') from zero in the fixed point iteration of (4.2). Accuracy is then quantified by computing the residuals $\|\Lambda' - F(\Lambda')\|_\infty$

$$F_i(\Lambda') \equiv \frac{\mu_i \theta_i^*}{1 - \mathcal{E}(\rho(\Lambda'))} \quad \forall \quad i = 1, \dots, N$$

The error is proportional to the residuals and is depicted in Fig. 4.4. The error values indicate that the variants provide a reasonable approximation to the Nash Equilibrium point. As opposed to the Newton-Raphson version, the logarithmic algorithm displays increasing accuracy with respect to number of users. This is due to the progressively better

log-linear approximation to the control-observation function as evident in Fig. 4.2. The increasing accuracy leads us to suspect that the logarithmic variant might asymptotically converge to the Nash equilibrium under certain limits.




Figure 4.4: Error vs Number of users: Logarithmic and Newton-Raphson variants.

4.4.7 Effect of Aggressiveness on System Convergence

User aggressiveness is characterized by β as defined in (4.5). It modifies the rate of user aggression as

$$\frac{d\alpha_i}{dk} = -\beta^k \log(\beta)$$

Hence as it approaches zero, players try to attain their optimal bandwidth more aggressively and in fewer time steps. The effect of aggression on system convergence for the three algorithms is shown in figs. 4.5, 4.6, 4.7. The number of iterations required for convergence increase as the variants become less aggressive, causing them to move slowly towards their equilibrium points.




Figure 4.5: Iterations vs β , Original algorithm.

4.4.8 Impact of User Demand on System Convergence

User demand is characterized using ϵ defined as $\epsilon \equiv 1 - \frac{\sum_{i=1}^N \theta_i^*}{K}$. As we vary ϵ from 1 to 0, user demand approaches the supply limit. We investigated the effect of ϵ on system convergence for the Logarithmic variant and $N = 2, 5$ and 10 users as illustrated in Fig. 4.8. The Newton-Raphson version also displayed analogous behavior. As the demand for bandwidth increase, it restricts the adaptability of each user. Any slight perturbation in demand due to noise becomes amplified slowing the convergence of the distributed algorithm.




Figure 4.6: Iterations vs β , Logarithmic variant.




Figure 4.7: Iterations vs β , Newton-Raphson variant.




Figure 4.8: Iterations vs ϵ , Logarithmic variant.

Chapter 5

Bi-level Optimization

The previous chapters mentioned in detail the application of the user rate game for bandwidth allocation in a single and multi-link circuit switched networks. We now take another look at the problem [32] and show that the distributed game paradigm can be viewed as a mapping from the vector of user desired bandwidths Θ^* to a feasible bandwidth allocation vector Θ^{eqm} . We then try to relate the noncooperative game to the problem involving the maximization of total user utilities. However, this user optimization is not done in isolation from the resource provider. Depending on the outputs received from its users, the network concurrently tries to solve its revenue maximization problem.

5.1 Rate Optimization

From a social welfare point of view, the ideal resource allocation among competing users should ensure that the total user utility is maximized. For the single link scenario of 3.1, the optimal bandwidth allocation is obtained by solving *Problem A*:

$$\begin{aligned} \max_{\Theta} \quad & \sum_{i=1}^N U_i(\theta_i) \\ & \sum_{i=1}^N \theta_i \leq C \end{aligned}$$

This is a constrained optimization problem with feasible vectors allocating bandwidth no greater than the available capacity C . Employing Lagrangian multipliers, we convert it into the following unconstrained problem:

$$\max_{\Theta, L} \sum_{i=1}^N U_i(\theta_i) - L(\sum_{i=1}^N \theta_i - C)$$

where L is the Lagrange multiplier for the constraint. The additive nature of the objective function indicates that this problem is separable. The linearity of the supply-demand constraint lends an economic interpretation to the Lagrangian L being the shadow price for the available bandwidth. The corresponding Karush-Kuhn-Tucker (KKT) conditions for optimality are:

$$\frac{\partial U_i(\theta_i)}{\partial \theta_i} = L \quad (5.1)$$

$$L \geq 0, \quad \sum_{i=1}^N \theta_i \leq C \quad (5.2)$$

$$L(\sum_{i=1}^N \theta_i - C) = 0. \quad (5.3)$$

If the optimal price L^\dagger is nonzero, by the above conditions it is also the market clearing price where all the capacity is utilized,

$$\sum_{i=1}^N \theta_i^\dagger = C \quad (5.4)$$

where θ_i^\dagger is the i^{th} user's optimal bandwidth. The resulting microeconomic love story - *supply meets demand* is illustrated in Fig. 5.1.

However the noncooperative setting precludes any computation of the objective function or its derivative as it requires knowledge about the utilities of individual users. We thus need to devise a distributed algorithm which can be used by each user to update her allocation without divulging utility information to other users or the network. The separability of the objective function suggests that individual users can instead try to solve the simpler one dimensional *problem B*

$$\max_{\theta_i} U_i(\theta_i) - L\theta_i \quad \forall i \in \mathcal{N}$$

This is the classic utility maximization problem of the individual user with L being a usage based bandwidth price.

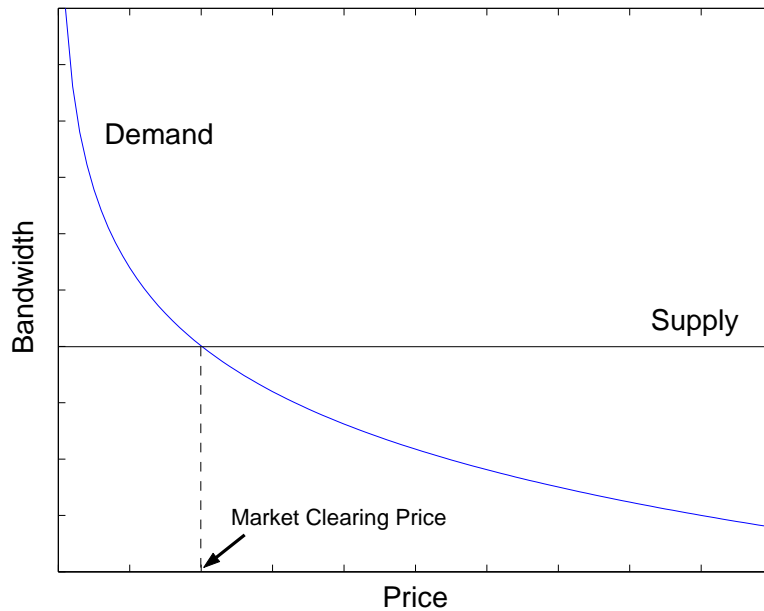


Figure 5.1: Bandwidth Supply and Demand

In the previous chapters, we had modeled user behavior using this profit maximization paradigm. Each user tried to maximize his/her net benefit which is the difference between the utility derived from bandwidth and the price paid to the network for its use. By appropriately choosing their call request rates, users then strived to attain their optimal bandwidths θ^* computed from (2.1). The strict concavity of the user utility functions ensures that the resulting optimal bandwidth vector Θ^* is unique. Although the objective function of *problem A* depends only on the bandwidth consumed by the user, we need to ensure that the user bandwidth vector Θ lies in the constraint set

$$S \equiv \{\Theta \in \mathfrak{R}^N \mid \sum_{i=1}^N \theta_i \leq C\}.$$

for the allocation to be feasible.

While playing the noncooperative rate game in (2.4), the user rates converge to a Nash equilibrium point Λ^{eqm} . The desired bandwidth vector Θ^* is thus mapped to an arrival vector Λ^{eqm} in the feasible rate set \mathcal{S} . However this mapping is only defined implicitly using (3.1). It may not be one-to-one and is dependent on the maximal rates. At the rate equilibrium, the corresponding bandwidth allocation vector is computed using

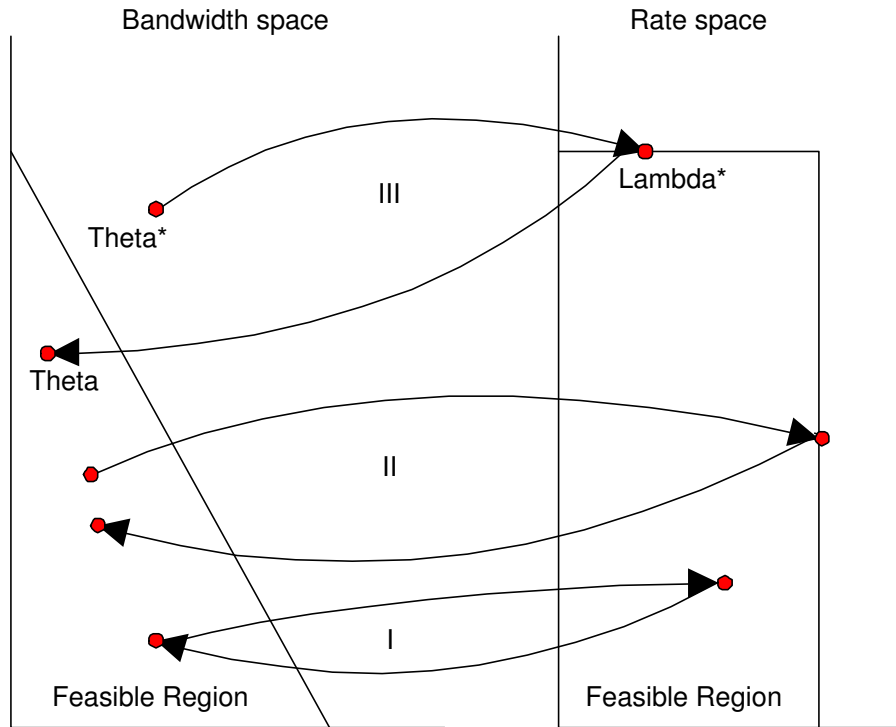


Figure 5.2: Mapping of Θ^* to feasible bandwidth region

(3.2). This mapping from Λ^{eqm} to Θ^{eqm} is however one-to-one. Further it maps the vector to the feasible region as proved by lemma 1 for the single link case and theorem 4 for the multi-link network. Thus the user game in effect maps the possibly infeasible optimal allocation vector to the feasible bandwidth allocation space.

Fig. 5.2 illustrates this mapping from Θ^* to Θ^{eqm} . Three different scenarios are possible during this feasibility transformation namely

- Case I: When the optimal vector is feasible and the maximal rates are large enough, the equilibrium and optimal bandwidth vectors become equal.

$$\lambda_r^{eqm} = \min\left\{\frac{\mu_r \theta_r^*}{1 - B_r(\Lambda^{eqm})}, \lambda_r^{max}\right\} = \frac{\mu_r \theta_r^*}{1 - B_r(\Lambda^{eqm})}$$

$$\Rightarrow \theta_r^{eqm} = \frac{\lambda_r^{eqm}(1 - B_r(\Lambda^{eqm}))}{\mu_r} = \theta_r^*$$

- Case II: When the optimal vector is feasible and the maximal rates are not large enough to include the Nash equilibrium point, the fixed point iterations end up at the

maximal rate boundary. Then the equilibrium and optimal bandwidth vectors will not be equal.

$$\begin{aligned} \exists r \quad s.t \quad \lambda_r^{eqm} &= \min\left\{\frac{\mu_r \theta_r^*}{1 - B_r(\Lambda^{eqm})}, \lambda_r^{max}\right\} = \lambda_r^{max} \\ \Rightarrow \theta_r^{eqm} &= \frac{\lambda_r^{eqm}(1 - B_r(\Lambda^{eqm}))}{\mu_r} = \frac{\lambda_r^{max}(1 - B_r(\Lambda^{eqm}))}{\mu_r} \neq \theta_r^* \end{aligned}$$

- Case III: When the optimal vector is not feasible, the Nash equilibrium always ends up at the maximal rate boundary. Analogous to case II, the equilibrium and optimal bandwidth vectors will then not be the same.

Let $(\Theta^\dagger, L^\dagger)$ be the optimal solution-Lagrange pair to *problem A*. Since they satisfy the optimality conditions (5.1), the allocation Θ^\dagger is feasible. When the maximal rates are not restrictive, case I in Fig. 5.2 applies leading to $\Theta^{eqm} = \Theta^*$. Thus the mapping does not disrupt the optimal solution to *problem A*. Similarly in *problem B* if the users are charged a positive market clearing price (5.4), the equilibrium allocation and bandwidth price also solve *problem A* due to their satisfying the KKT conditions (5.1).

5.2 Price Optimization

The emphasis in pricing literature has been on charging users a usage cost per bandwidth unit consumed, fixed at the marginal cost of production. While this is a Pareto efficient outcome which maximizes the benefit of all the users, it provides no venue for the network to recover its sunken costs. The paradox of an economically efficient outcome which is inimical to the provider's interests can be explained as follows. The efficiency criterion only mandates that the marginal unit be sold at marginal cost. Thus units produced before the marginal unit can be sold at a higher price and still maintain efficiency. But when all goods are sold at the same price, the provider is forced to sell them at the untenable marginal cost. The solution is thus to nudge the system to another Pareto efficient equilibrium using nonlinear and/or customer specific charging schemes. When the service provider is a profit maximizing entity, it would carry out price discrimination by charging a usage based fee such that the marginal willingness to pay equals the marginal cost. Here, marginal willingness to pay is defined as the willingness of a customer to pay for an incremental unit of the good. This implies that if a customer valued an additional unit of the good at more than

its production cost, she would be better off by purchasing the good. The provider will also profit from this transaction. Due to the law of diminishing marginal utility, the user's marginal willingness to pay will decrease as she accrues more and more goods. The network would then sell its last unit at a price equal to its marginal cost. At this equilibrium point, the marginal benefit accrued by both the producer and user is zero. This operating point also exhibits Pareto efficiency wherein the network garners all the social surplus.

We now demonstrate the relevance of differential pricing for advancing the network's objective. The network is a revenue maximizing entity with a revenue function

$$T(M, \Theta^{eqm}) \equiv \sum_{i=1}^N M_i \theta_i^{eqm},$$

which is a function of the market price and the equilibrium bandwidth allocated to the various users. The network chooses an appropriate market price to solve the optimization *Problem C*:

$$\begin{aligned} & \max_M T(M, \Theta^{eqm}) \\ & M \geq 0, A\Theta^{eqm} \leq C, \end{aligned}$$

namely maximizing its revenue while satisfying capacity constraints. For the single link scenario, we convert it into the following unconstrained problem using Lagrangian multipliers.

$$\max_{M, \alpha, \beta} M^T \Theta^{eqm} - \alpha (\sum_{i=1}^N \theta_i^{eqm} - C) + \beta^T M$$

The KKT conditions for *Problem C* can then be shown as:

$$\theta_r^{eqm} + M_r \frac{\partial \theta_r^{eqm}}{\partial M_r} - \alpha \frac{\partial \theta_r^{eqm}}{\partial M_r} + \beta_r = 0 \quad (5.5)$$

$$\alpha \geq 0, \sum_{i=1}^N \theta_i^{eqm} \leq C, \alpha (\sum_{i=1}^N \theta_i^{eqm} - C) = 0 \quad (5.6)$$

$$\beta \geq 0, M \geq 0, \beta^T M = 0 \quad (5.7)$$

The network would thus charge its users differently depending on their demand elasticity to price. The network cannot explicitly carry out this optimization as it is not cognizant of the dependence between the equilibrium allocation and the market price. The noncooperative users may loathe to part with this information or could even be unaware of it themselves. The scenario thus reduces to a bi-level optimization scheme.




Figure 5.3: Network-User interaction

Alternately the system could be viewed as a Stackelberg game [7, 8] with the network being the leader and the users acting as followers. Such a formulation has been studied to model networking routing and subsequent user behavior in [22]. The network initializes its algorithm by assigning an initial price $M(0)$ randomly or based on historical data. Note that while the initial conditions do not alter the final solution, they can significantly affect the *number of iterations* needed to reach the optimum. The users then treat the market price as given and compute the equilibrium bandwidth vector Θ^{eqm} by solving the user *problem B* and the bandwidth game. The network infers this allocation vector by observing the amount of bandwidth consumed, which is then employed to solve *problem C* with the updated Θ^{eqm} . Fig 5.3 illustrates this bi-level, network-user interaction.

In general the optimal market price M^* of *problem C* would not be the solution to the total utility maximization *problem A*. Since the network does not possess the derivative information for its objective function, we employ gradient-free algorithms to attain the optimal revenue. Stochastic approximation based schemes discussed in chapter 6 were chosen for this purpose owing to their excellent resilience to noisy measurements.

Chapter 6

Stochastic Approximation

We now present a brief introduction of the method of stochastic approximation (SA) for solving nonlinear root-finding problems in the presence of noisy measurements. Stochastic approximation is a stochastic analogue to the deterministic root finding techniques like steepest descent or Newton-Raphson algorithms. Often referred to as the Robbins-Monro algorithm, SA provides a general framework for analyzing the convergence properties of a wide variety of algorithms such as least-mean-squares (LMS), recursive-least-squares (RLS) and gradient-free SA. Root finding SA was introduced by Robbins and Monro in 1951 [40] with generalizations provided by Kiefer and Wolfowitz [21], Nevel'son and Has'minskii [30], Kushner and Clark [23].

Let X be the domain of allowable values for the decision variable x . We are interested in finding the values of $x \in X$ that are the roots of some vector-valued function $g(x)$. The root-finding function $g(x)$ is usually encountered while calculating the gradient for minimizing an objective function $T(x)$, i.e.,

$$g(x) = \frac{\partial T(x)}{\partial x},$$

Denote $\hat{x}(k)$ as the estimate for x at the k^{th} iteration of the algorithm under consideration. Noisy measurements of the gradient g at an estimate $\hat{x}(k)$ are represented as

$$Y(\hat{x}(k)) \equiv g(\hat{x}(k)) + e(\hat{x}(k)).$$

where e represents the noise term of dimension N . Observe that the noise is dependent on the state estimate $\hat{x}(k)$ indicating that the common assumption of independent and

identically distributed (i.i.d) noise will not apply since $\hat{x}(k)$ will be changing as the search process proceeds in time.

Our focus is in finding at least one root $x^* \in X^* \subseteq X \subseteq \mathfrak{R}^N$. Consider the steepest descent algorithm where

$$x(k+1) = x(k) - a(k)g(x(k)),$$

$a(k) > 0$ being the step size. An obvious stochastic extension would be to approximate $g(x(k))$ using multiple values of $Y(\hat{x}(k))$. Robbins and Monro showed this to be a wasteful use of the measurements since $g(\hat{x}(k))$ is merely an intermediate step in the process of attaining the desired x^* . Instead, they suggested a form of averaging across iterations which leads to a more effective use of the input information. We present below the unconstrained and constrained versions of the Robbins-Monro SA algorithm. Let a_k be a nonnegative “gain” value and $\Psi_x[\cdot]$ a mapping from a point outside X to a new point inside X .

Unconstrained:

$$\hat{x}(k+1) = \hat{x}(k) - a(k)Y(\hat{x}(k)) \tag{6.1}$$

Constrained:

$$\hat{x}(k+1) = \Psi_x[\hat{x}(k) - a(k)Y(\hat{x}(k))] \tag{6.2}$$

SA is thus derived from the steepest descent algorithm with $g(\hat{x}(k))$ being replaced with the noisy measurement $Y(\hat{x}(k))$. The contribution of SA however lies in the specification of the gain coefficients a_k to ensure proper across-iteration averaging and convergence to the root x^* . While these conditions may differ from those of its deterministic cousins, they also apply to steepest descent as it is a special case of SA.

6.1 Convergence of Stochastic Approximation

Most of the convergence results for SA are sufficient conditions and in the almost sure (a.s) sense. They have evolved out of two general settings - “statistics” conditions which impose statistical conditions on the gradient function $g(\cdot)$ and noise $e(\cdot)$ and “engineering” conditions based on Ordinary Differential Equation (ODE) theory. The latter is derived

from an ODE which roughly emulates the SA algorithm for large k as random effects disappear. The convergence properties of this deterministic differential equation are closely related to the convergence properties of (6.1). Neither of these conditions are special cases of each other and thus necessarily weaker. They apply when there is a unique root x^* to $g(\cdot)$ or when the objective function has local minima not different from the global minimum. The convergence conditions for the two settings have been taken from Spall [44].

6.1.1 “Statistics” conditions

The “statistics” conditions for the strong (a.s) convergence of the SA iterate $\hat{x}(k)$ are:

A.1 (Gain sequence) $a(k) > 0$, $a(k) \rightarrow 0$, $\sum_{k=0}^{\infty} a(k) = \infty$, and $\sum_{k=0}^{\infty} a(k)^2 < \infty$.

A.2 (Search direction) For some symmetric, positive definite matrix B and every $0 < \eta < 1$,

$$\inf_{\eta < \|x - x^*\| < 1/\eta} (x - x^*)^T Bg(x) > 0.$$

One may choose any convenient B such as the identity matrix \mathbf{I}_p .

A.3 (Mean-zero noise) $E[e(k, x)] = 0$ for all x and k .

A.4 (Growth and variance bounds) $\|g(x)\|^2 + E(\|e(k, x)\|^2) \leq c(1 + \|x\|^2)$ for all x and k and some $c > 0$.

These conditions are used for finding the appropriate convergence limits for the Finite Difference Stochastic Approximation in the next section. Condition A.1 is the most relevant since it indicates that the gain should be chosen with care so that its decay is neither too slow nor too fast. Thus as the iterate approaches the solution x^* , the gain should approach zero sufficiently fast ($a(k) \rightarrow 0$, $\sum_{k=0}^{\infty} a(k)^2 < \infty$) to damp out the noise effects. However it should also degrade at a sufficiently slow rate ($\sum_{k=0}^{\infty} a(k) = \infty$) to avoid premature convergence of the algorithm. Given the desirability for a gain sequence that balances algorithm stability in the earlier iterations with nonnegligible step sizes in the later iterations, a recommended gain form is

$$a(k) = \frac{a}{(k + 1 + A)^\alpha}, \tag{6.3}$$

where $A \geq 0$ is the stability constant. When $A = 0$ and a large a is chosen to induce large steps in the later iterates, it may lead to unstable behavior in the beginning. However if a is chosen to be small, performance in the latter cases is sacrificed for robust behavior at the onset of SA. Thus a positive value of A is often preferred. For values of $0.5 < \alpha \leq 1$, A is chosen to be approximately 5 to 10 percent of the total number of expected or allowed iterations.

6.1.2 “Engineering” conditions

For the ODE approach mentioned above, conditions for the “engineering” type are given below. These are used for finding the appropriate convergence limits for the Simultaneous Perturbation Stochastic Approximation (SPSA) in the next section.

B.1 (Gain sequence) $a_k > 0, a_k \rightarrow 0, \sum_{k=0}^{\infty} a_k = \infty$.

B.2 (Relationship to ODE) Let $g(x)$ be continuous on \mathfrak{R}^p . With $Z(t) \in \mathfrak{R}^N$ representing a time-varying function, suppose that the differential equation given by

$$\frac{dZ(t)}{dt} = -g(Z(t))$$

has an asymptotically stable equilibrium point at x^* .

B.3 (Iterate boundedness) $\sup_{k \geq 0} \|\hat{x}(k)\| < \infty$ a.s. Further $\hat{x}(k)$ lies in a compact subset of the domain of attraction for the differential equation in B.2 infinitely often.

B.4 (Bounded variance property of measurement error) Let $\mathfrak{F}_k \equiv \{\hat{x}(0), \dots, \hat{x}(k)\}$. Let the bias for the k^{th} iterate be $b(k) = E[e(k, \hat{x}(k)) | \mathfrak{F}_k]$. Then $E[\|\sum_{k=0}^{\infty} a(k)(e(k) - b(k))\|^2] < \infty$.

B.5 (Disappearing bias) $\sup_{k \geq 0} \|b(k)\| < \infty$ a.s and $b(k) \rightarrow 0$ a.s as $k \rightarrow \infty$.

Unlike the “statistics” conditions, it is assumed that the gradient function $g(x)$ is continuous. The boundedness condition B.3 has been somewhat controversial in the literature as it imposes a requirement on the very iterate that one is trying to analyze. Condition B.4 ensures that the martingale convergence theorem from probability theory can be used to cope up with the noise effects in the SA recursion. B.5 generalizes the mean-zero noise restriction imposed by A.3 by considering noise functions whose mean converges

to zero as well. This extension is essential for analyzing the convergence of gradient free SA methods like SPSA.

6.2 Gradient-Free Algorithms

In our bi-level optimization strategy, the exchange of information between users and the network follows the schema of Fig. 5.3. Thus, for a particular market price M , the network can only observe the corresponding bandwidth allocation $\Theta(M)$ among the users. Since the utility information of the users are considered private, the network provider is unable to obtain the higher derivatives $(\partial^n \theta_i^{eqm} / \partial M_i^n)$ of the bandwidth allocated with respect to the market price. Only the revenue measurements $y(M) = T(M, \Theta^{eqm}(M)) + \epsilon(M)$ are available for various values of the market price M . We thus resort to gradient-free algorithms which are based on an approximation to the gradient formed from noisy measurements of the objective function. Note that the notion of “gradient-free” is applicable only to the implementation of the algorithm. It does not refer to the absence of a gradient, which is required to guarantee convergence.

6.2.1 Finite Difference SA (FDSA)

Based on the seminal work by Robbins and Monro, an SA algorithm based on the Finite Difference (FD) gradient approximation was introduced for the scalar and multivariate cases by Kiefer-Wolfowitz and Blum respectively. Gradient approximation is carried out by the finite difference method using small one-at-a-time changes on each of the individual elements of M . The possibly noisy value of the network revenue function $T(M)$ (see Section 5.2) is then measured after each change. Once the measurements of the objective function have been collected for perturbations in each element of M , the gradient approximation may be calculated. This is motivated directly from the definition of the gradient as a collection of derivatives for each of the components in M , holding all other components fixed.

$$g_i(M) = \frac{\partial T}{\partial M_i} = \lim_{\Delta \rightarrow 0} \frac{T(M_1, \dots, M_{i-1}, M_i + \Delta, M_{i+1}, \dots, M_N) - T(M)}{\Delta}.$$

The network tries to solve its revenue maximization *Problem C* by recursively

updating its market prices in the general SA form:

$$\hat{M}(k+1) = \hat{M}(k) - a(k) \hat{g}(k, \hat{M}(k)), \quad (6.4)$$

where $\hat{g}(k, \hat{M}(k))$ is the estimate of the gradient ∇T at the iterate $\hat{M}(k)$ based on measurements of the revenue function. Thus (6.4) is analogous to the basic Robbins-Monro algorithm (6.1) with the gradient estimate $\hat{g}(k, \hat{M}(k))$ replacing the gradient $g(k, \hat{M}(k))$. The crucial component of (6.4) is the gradient approximation which may be obtained from the measurements of $y(\hat{M}(k))$ and $y(\hat{M}(k)+\text{perturbation})$ or from two measurements of $y(\hat{M}(k)\pm \text{perturbation})$, called one-sided and two-sided approximations, respectively. We employ the latter to approximate the gradient as

$$\hat{g}(k, \hat{M}(k)) = \begin{bmatrix} \frac{y(\hat{M}(k)+c(k)\eta_1) - y(\hat{M}(k)-c(k)\eta_1)}{2c(k)} \\ \vdots \\ \frac{y(\hat{M}(k)+c(k)\eta_N) - y(\hat{M}(k)-c(k)\eta_N)}{2c(k)} \end{bmatrix}, \quad (6.5)$$

where η_i denotes a vector with a 1 in the i th place and 0's elsewhere and $c(k) > 0$ defines the difference magnitude. The pairs $\{a(k), c(k)\}$ are the gain sequences for the FDSA algorithm. Thus the gradient approximation is formed by perturbing the components of the market price $\hat{M}(k)$ one a time and collecting a revenue measurement at each of the perturbations. This requires a total of $2N$ measurements for a two sided gradient approximation.

This algorithm relies only on revenue computations to optimize the network objective. In fact, we need to find only the difference between the two values of the revenue function. However, this method would be costly for higher dimensional problems since a measurement has to be obtained for each element of M . Because a typical implementation for optimization requires several iterations with approximations at each step, this measurement cost could be prohibitive. Every time the network declares a new set of bandwidth prices, the users are forced to play a bandwidth game which eventually converges to a Nash equilibrium. Chapter 4 indicates that the number of iterations required for convergence will increase as measurement noise and user population goes up.

The convergence theory of the FDSA algorithm is similar to that of the root finding SA detailed in Section 6.1. However an additional difficulty arises due to the bias in $\hat{g}(k, \hat{M}(k))$ as an estimator of $g(k, \hat{M}(k))$ and the need to control the extra gain sequence $c(k)$. This is because the ‘‘statistics’’ conditions (A.1-A.4) are based on unbiased estimates of $g(\cdot)$ for all k . The introduction of the estimate leads to an additional bias of $O(c_k^2)$ in

the computation of $E[g(k, \hat{M}(k)) | \mathfrak{S}_k]$. Further $\hat{g}(k, \hat{M}(k))$ is an unbiased estimator of the gradient in the special case where the objective function is quadratic.

Some guidelines have been provided for specific choices of the gain sequences namely,

$$a(k) = \frac{a}{(k+1+A)^\alpha}, \quad c(k) = \frac{c}{(k+1)^\gamma} \quad (6.6)$$

where a, c, α and γ are strictly positive. Practical values of α and γ that are effectively as low as possible while satisfying A.1' are 0.602 and 0.101, respectively. To cope with noise effects, c is set at a level approximately equal to the standard deviation of the measurement noise. This helps keep the elements of the gradient estimate from growing excessively large before $a(k)$ stabilizes the search process. A is then chosen to be around 10 percent of the maximum number of iterations. We then choose the most critical coefficient a . In making this choice, we initially picked $a_{temp,i}, i = 1, \dots, N$ such that $[a_i^{temp}/(A+1)^{0.602}]$ when multiplied with the corresponding element of the initial gradient $g(0)$ is approximately equal to the desired change in magnitude of the M_i in the early iterations. We then repeated this process for several replications and chose a as the minimum of all a_i^{temp} .

6.2.2 Simultaneous Perturbation SA (SPSA)

Compared to FDSA, the main benefit of Simultaneous Perturbation Stochastic Approximation is a reduction in the number of loss measurements required to achieve a given level of accuracy in the optimization process. SPSA has found applications in areas as varied as queuing systems, simulation-based optimization, aircraft design and neural network training. The basic algorithm (6.4) is augmented with an innovative method to compute the gradient approximation. While a two sided FDSA requires a total of $2N$ measurements per iteration, SPSA computes the gradient using *just two measurements*. The two-sided simultaneous perturbation gradient approximation is obtained by randomly perturbing all the elements of $\hat{M}(k)$ leading to

$$\hat{g}(k, \hat{M}(k)) = \begin{bmatrix} \frac{y(\hat{M}(k)+c(k)\Delta_k) - y(\hat{M}(k)-c(k)\Delta_k)}{2c(k)\Delta_{k1}} \\ \vdots \\ \frac{y(\hat{M}(k)+c(k)\Delta_k) - y(\hat{M}(k)-c(k)\Delta_k)}{2c(k)\Delta_{kN}} \end{bmatrix}, \quad (6.7)$$

where the mean-zero N -dimensional random perturbation vector $\Delta_k = [\Delta_{k1}, \dots, \Delta_{kN}]^T$ has a symmetric Bernoulli ± 1 distribution and $c(k)$ is a positive scalar.

Since the numerator is the same in all N components of $\hat{g}(k, \hat{M}(k))$, only 2 loss measurements are needed to estimate the gradient in SPSA regardless of the number of users N . This measurement savings per iteration provides a potential way to achieve considerable advantage over FDSA in the total number of measurements required to estimate M when p is large. This potential is realized if the number of iterations required for effective convergence to an optimum M^* does not increase in a way to negate the measurement savings per gradient iteration at each iteration. It is clear that the FD approximation (6.4) will generally be superior to that obtained through (6.7). Under reasonably general conditions, the SPSA and FDSA recursions attain the same level of statistical accuracy for a given number of iterations even though SPSA uses only $1/N$ times the number of function evaluations of FDSA.

$$\frac{\text{number of } y(\cdot) \text{ values in SPSA}}{\text{number of } y(\cdot) \text{ values in FDSA}} \rightarrow \frac{1}{N} \quad (6.8)$$

as the number of loss measurements in both procedures gets large. This implies that the N -fold savings per iteration for gradient approximation translates directly into a N -fold savings in the overall optimization process. Note that this is under the assumption that both FDSA and SPSA use the same gain sequences $a(k)$ and $c(k)$.

6.3 SPSA Pricing algorithm

In this section we present a step-by-step summary of a pricing algorithm based on SPSA.

Step 0 (Initialization and coefficient selection) Initialize iteration index $k = 0$. Pick initial price $\hat{M}(0)$ and set $\alpha = 0.602, \gamma = 0.101$. The coefficients a, A and c are determined by guidelines in Section 6.1.

Step 1 (Generation of the simultaneous perturbation vector) Generate by Monte Carlo a p -dimensional perturbation vector Δ_k , where each of the p components of Δ_k are independently generated from a Bernoulli ± 1 distribution with probability of $1/2$ for each outcome.

Step 2 (Objective function evaluations) Obtain two measurements of the objective function based on the simultaneous perturbation around the current price $\hat{M}(k)$.

Step 3 (Gradient approximation) Generate the simultaneous perturbation approximation to the unknown gradient $g(\hat{M}(k))$ according to (6.7). When the noise effects $\epsilon(k)$ are relatively large, several gradient approximations computed at $\hat{M}(k)$ could be averaged.

Step 4 (Price updating) Use (6.2) to update the network price $\hat{M}(k)$ to a new value $\hat{M}(k+1)$. The mapping function $\Psi_M(\cdot) \equiv \max(\cdot, \mathbf{0})$ is employed to maintain feasibility. The SA formula (6.2) was designed to minimize the objective function. Since our objective is revenue maximization, we update the prices as follows

$$\hat{M}(k+1) = \max(\hat{M}(k) + a(k)\hat{g}(k, \hat{M}(k)), \mathbf{0}).$$

Step 5 (Iteration or Termination) Replace k with $k+1$ and go to step 1. Terminate the algorithm if there is little change in several successive iterates or if the maximum allowable number of iterations has been reached.

The choice of the gain sequences is critical to the performance. With α, γ specified in step 0 of the algorithm above, it is found that in a high-noise setting it is necessary to pick a smaller a and larger c than those employed in a low-noise setting.

6.4 Numerical Results

We ran numerical experiments for the perfect information regime. All users were updated with the latest rate vector while adapting their respective individual call request rates. In this case due to the presence of noise-free measurements, FDSA is the same as the classic Finite Difference gradient-free algorithm. However the corresponding sample paths generated by SPSA would not be the same due to the influence of random perturbations used in gradient approximation. We ran 50 runs to obtain statistically significant results. The 95% confidence intervals (C.I) were calculated using

$$\text{C.I} = \bar{M} \pm z_{0.025} \frac{s}{\sqrt{n}} = \bar{M} \pm 1.96 \frac{s}{\sqrt{n}} \quad (6.9)$$

where \bar{M} is the sample average, $z_{0.025}$ is the, s is the sample standard deviation and n is the total number of runs.




Figure 6.1: Comparison of efficiency for FDSA and SPSA.

6.4.1 Efficiency of SPSA

The efficiency and scalability of algorithms are key issues in solving any problem of interest. Performance may be measured in terms of computer run time, number of algorithm iterations, and number of objective function evaluations. We employ the latter criterion to compare the FDSA and SPSA algorithms. This is because computer run times are dependent on the specific machine employed, operating system in place and the processes running in memory while the simulation is in motion. Thus the results may not readily transfer from one setting to another. Iterations can also be misleading, as some algorithms may have iterations which are much more computationally intensive than others. Hence an algorithm converging in fewer iterations may not have lesser computational cost than others which use more iterations but with lower per-iteration demands. Thus we compare the performance of our algorithms with respect to the number of revenue evaluations. The motivation behind this approach is that in practice the $T(\cdot)$ measurements are usually the dominant cost in the optimization process especially as the number of users N becomes large.

Fig. 6.1 compares the two algorithms while coordinating a single link network comprising of 10 identical users and 100 circuits. Note that FDSA and SPSA compare favorably on an iteration-by-iteration basis. However, SPSA uses only two revenue measurements per iteration while FDSA uses 20. This 10-fold average savings per iteration leads to the large savings in total measurements for the full number of iterations. Due to the superior performance of SPSA, we henceforth employed it as our primary price adjustment algorithm.

6.4.2 Impact of Number of Users on Revenue Generation




Figure 6.2: Revenue generation for varying number of users using FDSA.

The network adjusted its prices for a single link comprising of 100 circuits and users ranging from 2, 5 and 10. The users had identical utility functions with coefficient b set as 10. From the KKT conditions (5.5) for the revenue maximization problem, we obtain

$$N\left(\frac{b}{M^*} - 1\right) = 100.$$

Thus the optimal market revenue turns out to be

$$T(M^*) = NM\left(\frac{b}{M} - 1\right) = 100\frac{bN}{100 + N},$$




Figure 6.3: Revenue generation for varying number of users using SPSA.

which are tabulated below. In the case of a 10 user system, the market clearing effects near the optimal revenue are illustrated in Fig. 6.4.

| No. of Users | $T(M^*)$ |
|--------------|----------|
| 2 | 19.61 |
| 5 | 47.62 |
| 10 | 90.91 |

Fig. 6.2 and 6.3 illustrate the network revenue for varying iterations and number of users. In both cases, we observe that the network objective approaches the optimal values tabulated above. The earnings accrued by the network can be seen to increase with increasing number of users. This is because a larger population of users leads to a larger demand for resources which encourages the provider to sell her goods at a higher price.

6.4.3 Effect of Initial Price on Convergence

Our numerical studies indicated that the SPSA algorithm converges to an equilibrium irrespective of the initial price vector $M(0)$. However, the starting price is a critical factor in the convergence to the equilibrium network price. Fig. 6.5 depicts these facts in the case of two identical users sharing a single link.

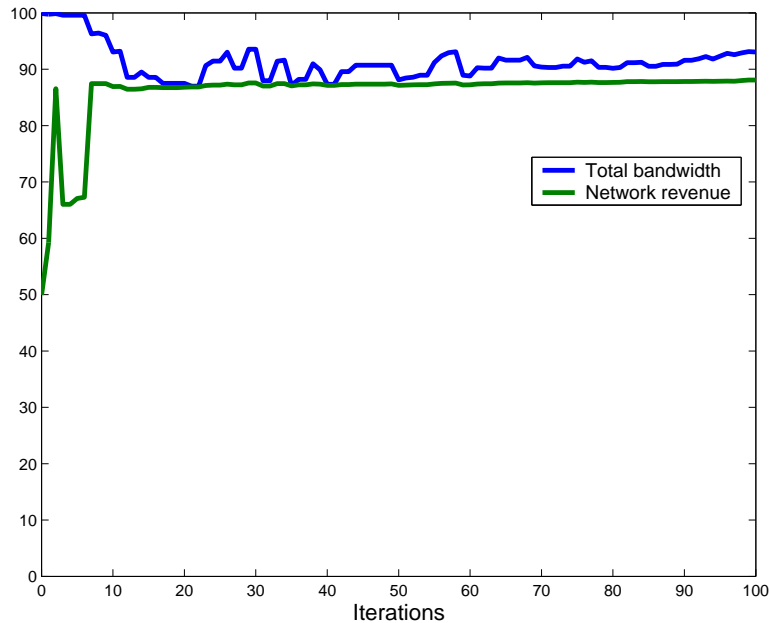


Figure 6.4: Market clearing effects for a system of 10 users.

6.4.4 Triangle Network

We now turn our attention to the triangle network of Fig. 3.11. Users 1 and 2 with utility coefficients $b_1 = 10, b_2 = 50$ demand 9 and 49 circuits from the network respectively. A link capacity vector $C = [10 \ 10 \ 10]^T$ was chosen to simulate a bottleneck in link 3. Unlike the scenario in 3.4.3, both the network and users were involved in this bilevel adaptation scheme. The system eventually settled down to an equilibrium with the capacity of link 3 allocated as per Fig. 6.6. The resource provider maximized her revenue by employing price discrimination. User 2 with its higher demand was willing to pay more for the resource than user 1. The network thus charged her 43% more than user 1 (see Fig. 6.7) and allocated 83% of link 3 capacity.

In the presence of competition for scarce resources, users may try to modify their parameters to capture a large portion of the available capacity. In our triangle network, the bone of contention is the bottle neck link 3. In our numerical experiment, user 2 modified her maximum arrival rate λ_2^{max} while user 1 maintained a static maximal rate. The aspect ratio $\lambda_2^{max}/\lambda_1^{max}$ was varied from 1 to 10. We maintained the market price at a constant $M = [1 \ 1]^T$. Since the bandwidth price is lower than the market clearing price,




Figure 6.5: Effect of Initial Price on Convergence.

user demands outstrip the available link capacity. The equilibrium rates thus settle at the boundary of the maximal rate box. As λ_2^{max} is increased, the component of user 2 in the final arrival rate dwarfs that of user 1. This in turn leads to a higher bandwidth allocation in favor of the second user. Fig. 6.8 confirms this analysis with respect to the bottle neck link. User 1 gets a better deal when her maximal rate is raised from 10 to 100. This has serious implications for fairness in real world situations. The maximal arrival rates could be an inherent property of computing power of the network processor and transmission capacity. Hence, by boosting its maximum rate, any entity with superior processing power and technology could gain an upper hand over its less fortunate peers.




Figure 6.6: Capacity allocation in the Bottleneck Link.




Figure 6.7: Price and Revenue Variation in the Bottleneck Scenario.




Figure 6.8: Sensitivity of bandwidth allocation to maximal rates in the bottleneck link.

Chapter 7

Conclusion

This dissertation has presented a unified pricing model for regulating the interaction between a connection oriented network and its users. This chapter summarizes the work and suggests extensions for future work.

7.1 Summary

We envisage the future Internet to be composed of MPLS/GMPLS networks running over optical networks enabled with Dense Wavelength Division Multiplexing technology. The implications of such a scenario are two-fold. Firstly, they portend the resurgence of connection oriented networks albeit in conjunction with IP networks. Secondly, their Gigabit speeds would foster the development of bandwidth intensive applications like Video-on-Demand and online, 3D gaming. These would require a higher quality of service different from the current best effort service model. To regulate such QoS-aware applications, network providers will have to introduce class sensitive pricing into the Internet. In this work, we have tried to provide an answer to the problems posed above.

We developed an integrated model to incorporate the interaction between users and the network in a connection oriented setting. A noncooperative community of users that shared a network was studied using the tools of game theory. The resultant Nash equilibrium was studied for the single link and multi-link cases. The convergence to the fixed

point and the impact of maximal rates on the equilibrium were also studied. Upper bound based variants helped speed up the computations in exchange for higher equilibrium rates. The user game was then extended to encompass the imperfect information regime where users adapted their rates based on feedback from the system. Three algorithms based on Recursive Least Squares were proposed and their convergence, accuracy and scalability was validated by simulation results. We then obtained a different perspective of the rate game where it was portrayed as a mapping to impose feasibility. This was then tied together with a bi-level optimization model where the network is the leader and the users are followers in a Stackelberg game. Novel gradient free Stochastic Approximation algorithms for revenue maximization were proposed and validated for the single link and triangle network. Thus a comprehensive suite of user and network models and algorithms were developed, tested and validated for varying price and network scenarios.

7.2 Future Work

A possible avenue for future work would be the extension of algorithms developed in Chapter 4 to operate in a general, multi-link network. However the route blocking probability experienced by any user is a complicated function of the blocking probabilities of other users as well as the network capacity. Thus for any feedback based algorithm to be stable, the observability of the state space has to be enhanced. This could be possible only if the users exchange their rate information or blocking information to other users. Such an arrangement may lead to an erosion of privacy as well as additional signalling burden and delay.

Further more the linear RLS algorithms could be replaced by their nonlinear counterparts like stochastic gradient based techniques. Preliminary forays in this direction were not successful as they were unable to outperform the logarithmic RLS variant for the single link case. Still, an algorithm suite encompassing an general network with arbitrary number of links and users would be a significant achievement in the field of telecommunications. The robustness of such schemes in the presence of artificially injected randomness offers scope for further study. As the randomness in the measurement data increase, the distributed optimization algorithms may push the system into a chaotic orbit.

In the real world, the network objective function $T(M, \Theta)$ could often represent

the total profit garnered. The network would then have to price its resources at a value higher than the cost to produce them. It would be interesting to see how the profit maximization and revenue maximization cases differ in terms of the bandwidth allocation. In this dissertation, we have only considered users with logarithmic utility functions. Ideally user utilities should be empirically calculated from real world demand data.

The algorithms in this dissertation could be incorporated into intelligent agents which could be tested in a controlled lab experiment. Human users would allocate a daily budget and the agent would have to decide on a QoS vector to optimize the net benefit under a rolling horizon. They should thus have the ability to interface with users and exchange messages over the network with other agents possibly using Session Initiation Protocol.

Another dimension of the agent negotiation approach discussed in Chapter 2 is the effect of asynchrony on system performance. When rate information is transmitted over the network, delays are inevitable due to packet losses, processing and propagation. The fixed point computation in 2.3 should be robust to such delays to ensure fast convergence to the Nash equilibrium. Each agent could follow a variety of approaches to adapt to information delays. It could wait for all the updated information ($N - 1$ updates) to arrive prior to processing or alternatively let the computation proceed as and when any new updates occur. Such asynchronous computations [46] are commonly implemented by operating systems and microprocessors for deciding new allocation and scheduling policies.

Bibliography

- [1] S. Alouf, E. Altman, and P. Nain. Optimal on-line estimation of the size of a dynamic multicast group. In *Proc. of INFOCOM 2002*, volume 2, pages 1109–1118, June 2002.
- [2] Tansu Alpcan and Tamer Basar. A utility-based congestion control scheme for Internet-style networks with delay. In *Proc. of INFOCOM 2003*, June 2003.
- [3] Aristos Aresti, Bobby M. Ninan, and Mihail Devetsikiotis. Resource allocation games in connection-oriented networks under imperfect information. To appear in *Proc. ICC*, 2004.
- [4] K. C. Border. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, London, 1985.
- [5] Ron Cocchi, Scott Shenker, Deborah Estrin, and Lixia Zhang. Pricing in computer networks: motivation, formulation, and example. *IEEE/ACM Transactions on Networking*, 1(6):614–627, 1993.
- [6] Carlos de Paz Martinez. A nonlinear dynamics analysis of overflow traffic. *Comunicaciones De Telefonica I+D*, (19):93–100, December 2000.
- [7] C. Douligeris and R. Mazumdar. Multilevel flow control of queues. In *Johns Hopkins Conference on Information Sciences*, Baltimore, MD, March 1989.
- [8] A.A. Economides and J.A. Silvester. Priority load sharing: an approach using Stackelberg games. In *Proc. Allerton Conference on Communication, Control, and Computing*, 1990.

- [9] R. Edell and P. Varaiya. Providing internet access: What we learn from INDEX. *IEEE Network*, 13(5), Sept-Oct 1999.
- [10] M. Falkner, M. Devetsikiotis, and I. Lambadaris. An overview of pricing concepts for broadband IP networks. *IEEE Communications Surveys*, 3(2), 2000.
- [11] Andras Farago. Blocking probability estimation for general traffic under incomplete information. In *Proc. ICC'2000*, New Orleans, May 2000.
- [12] Sally Floyd. TCP and Explicit Congestion Notification. *ACM Computer Communications Review*, 24:10–23, 1994.
- [13] R. Gibbens and F. Kelly. Resource pricing and the evolution of congestion control, 1998.
- [14] E. J. Hernandez-Valencia, L. Benmohamed, S. Chong, and R. Nagarajan. Rate control algorithms for the ATM ABR service. *Europ. Trans. Telecom*, 8(1), 1997.
- [15] P. J. Hunt. *Limit Theorems for Stochastic Loss Networks*. PhD thesis, University of Cambridge, 1990.
- [16] A. A. Jagers and E. A. van Doorn. On the continued Erlang loss function. *Operations Research Letters*, 5:43–46, 1986.
- [17] K. Kar, S. Sarkar, and L. Tassiulas. A simple rate control algorithm for maximizing total user utility. In *Proc. INFOCOM*, 2001.
- [18] F. P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [19] F.P Kelly, A.K. Maulloo, and D.K.H. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [20] Frank P. Kelly. Blocking in large circuit-switched networks. *Advances in Applied Probability*, 1986.
- [21] J. Kiefer and J. Wolfowitz. Stochastic estimation of a regression function. *Annals of Mathematical Statistics*, 23:462–466, 1952.

- [22] Y.A. Korilis, A.A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Trans. Networking*, 5(1), 1997.
- [23] H. J. Kushner and D. S. Clark. *Stochastic Approximation Methods for Constrained and Unconstrained Systems*. Springer-Verlag, New York, 1978.
- [24] R. La and V. Anantharam. Charge sensitive TCP and rate control in the internet. In *Proc. IEEE INFOCOM*, 2000.
- [25] Will E. Leland, Murad S. Taqq, Walter Willinger, and Daniel V. Wilson. On the self-similar nature of Ethernet traffic. In Deepinder P. Sidhu, editor, *ACM SIGCOMM*, pages 183–193, San Francisco, California, 1993.
- [26] Steven H. Low and David E. Lapsley. Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, 1999.
- [27] J. K. Mackie-Mason and H. R. Varian. Pricing the Internet. In *International Conference on Telecommunication Systems Modeling*, pages 378–393, March 1994.
- [28] T. Magedanz, K. Rothermel, and S. Krause. Intelligent agents: An emerging technology for next generation telecommunications? In *INFOCOM*, San Francisco, CA, USA, 24–28 1996.
- [29] E. J. Messerli. Proof of a convexity property of the Erlang B formula. *The Bell System Technical Journal*, 51:951–953, 1972.
- [30] M. B. Nevel’son and R. Z. Has’minskii. *Stochastic Approximation and Recursive Estimation*. American Mathematical Society, Providence, RI, 1973.
- [31] B. M. Ninan and M. Devetsikiotis. Pricing mediated bandwidth allocation for the next generation Internet. In *Proc. GLOBECOM*, 2003.
- [32] B. M. Ninan and M. Devetsikiotis. *Game-Theoretic Resource Pricing for the Next Generation Internet*, volume 25 of *Performance Evaluation and Planning Methods for the Next Generation Internet*. Kluwer, 2005.
- [33] B. M. Ninan, G. Kesidis, and M. Devetsikiotis. A simulation study of non-cooperative pricing strategies for circuit-switched optical networks. In *Proc. ACM/IEEE MAS-COTS*, 2002.

- [34] A. M. Odlyzko. Paris metro pricing for the internet. *Proc. ACM Conference on Electronic Commerce*, pages 140–147, 1999.
- [35] Andrew M. Odlyzko. Internet pricing and the history of communications. *Computer Networks*, 36:493–517, 2001.
- [36] Andrew M. Odlyzko. Internet traffic growth: Sources and implications. In W. Weierhausen, A. K. Dutta, and K. I. Sato, editors, *Optical Transmission Systems and Equipment for WDM networking II*, volume 5247, pages 1–15, 2003.
- [37] Andrew M. Odlyzko. The many paradoxes of broadband. *First Monday*, 8(9):1–15, September 2003.
- [38] Andrew M. Odlyzko. Pricing and architecture of the Internet: Historical perspectives from telecommunications and transportation. January 2004.
- [39] Reza Rejaie, Mark Handley, and Deborah Estrin. RAP: An end-to-end rate-based congestion control mechanism for realtime streams in the internet. In *INFOCOM*, pages 1337–1345, 1999.
- [40] H. Robbins and S. Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22:400–407, 1951.
- [41] C. U. Saraydar, N.B. Mandayam, and D.J. Goodman. Power control in a multicell CDMA data system using pricing. *IEEE VTC*, pages 484–491, 2000.
- [42] S. Shenker. Fundamental design issues for the future Internet. *IEEE JSAC*, 13:1176–1188, 1995.
- [43] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review*, 26:19–43, April 1996.
- [44] James C. Spall. *Introduction to Stochastic Search and Optimization*. Discrete Mathematics and Optimization. Wiley-Interscience, 2003.
- [45] Robert F. Stengel. *Optimal Control and Estimation*. Dover Publications, 1994.

- [46] J.N. Tsitsiklis, D.P. Bertsekas, and M. Athans. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. *IEEE Transactions on Automatic Control*, 31(9):803–812, 1986.
- [47] Hal R. Varian. Differential pricing and efficiency. *First Monday*, 1(2), August 1996.
- [48] R. W. Wolff. *Stochastic Modeling and Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [49] X. Xiao and L. M. Ni. Internet QoS: A big picture. *IEEE Network*, 13(2):8–18, March 1999.
- [50] H. Yaiche, R. R. Mazumdar, and C. Rosenberg. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, 2000.