

Title	The Microsoft Congestion Pricing and Network Games Bibliography
Creators	Dukes, W.M.B.
Date	1999
Citation	Dukes, W.M.B. (1999) The Microsoft Congestion Pricing and Network Games Bibliography. (Preprint)
URL	https://dair.dias.ie/id/eprint/430/
DOI	DIAS-STP-00-10

CONGESTION PRICING AND NETWORK GAMES BIBLIOGRAPHY

W.M.B. Dukes
School of Theoretical Physics
Dublin Institute for Advanced Studies
10 Burlington Rd.
Dublin 4
Ireland

DIAS-STP-00-10

December, 1999

Abstract

We review the core papers related to the Microsoft Research Project “Congestion Pricing and a Distributed Game”. The papers have been selected from several different groups who have considered different aspects of the network congestion and routing problems along with some material on the economics of pricing schemes for proposed networks.

Contents

1	Introduction	2
2	The Abstracts	4
2.1	The Cambridge Papers	4
	Charging and rate control for elastic traffic	4
	Rate control in communication networks: shadow prices, proportional fairness and stability	5
	Resource pricing and the evolution of congestion control	6
	Differential QoS and Pricing in Networks: where flow-control meets game theory	6
	Congestion Pricing for Congestion Avoidance	7
	Mathematical modelling of the Internet	8
	An Intelligent Agent for Optimizing QoS-for-Money in Priced ABR Connections	8
	Optimization Flow Control, I: Basic Algorithm and Convergence	9
	Bandwidth sharing: objectives and algorithms	10
	A modest proposal for preventing Internet congestion	11
2.2	Routing	13
	Fairness in Network Optimal Flow Control: Optimality of Product Forms . . .	13
	A Game Theoretic Perspective to Flow Control in Telecommunication Networks	13
	Competitive Routing in Multi-User Communication Networks	13
	Achieving Network Optima Using Stackelberg Routing Strategies	14
	Capacity Allocation under Non-cooperative Routing	15
	Optimal Routing Control: Game Theoretic Approach	15
2.3	Pricing	17
	Pricing the Internet	17
	Some Economics of the Internet	17
	Adding service discrimination to the internet	18

1 Introduction

“Necessity,” says the proverb, “is the mother of invention.” The past decade has seen a huge increase in the use of the Internet, along with a similar increase of the problems which come with it. The most annoying problem experienced by users is congestion. The current transmission control protocol(TCP) is not longer well suited to all the traffic types which are being transferred these days and the ability for some users to implement even greedier TCP algorithms on their end-systems to obtain even more bandwidth is causing serious concern.

The reason for this is clearly the zero-cost policy of the Internet and as long as this is the case, congestion will continue to be a problem. It is thought that future networks may be priced in some way and we review some papers investigating this idea of network economics.

In the presence of any pricing scheme, users will have tradeoffs between what their objective is, the cost of such an objective and how valuable this objective is to the user relative to the cost it incurs. This provides the link between the theory of network flows and game theory. Different users having different objectives which affect one another. The papers we look at in this bibliography exhibit such a link.

The models considered differ from paper to paper but each has its own perspective on what aspects of the network it seeks to examine. Flow control problems have been considered by Douligeris & Mazumdar[DM92], Bovopoulos & Lazar[BL87] and Hsiao & Lazar[Laz91]. Lee & Cohen[LC85] study customer allocation in a system of parallel M/M/c queues. Dziong & Mason[DM91] study a call admission control problem. Routing problems are studied in La & Anantharam[LA98], Orda,Rom & Shimkin[AOS93], Economides & Silvester[ES91] and Yamaoka & Sakai[YS95] [YS96]. Lazar,Orda & Pendarakis[ALP95] investigate assigning bandwidth to different virtual paths. Shenker[She90] investigates an internetwork gateway problem. Traffic from end-systems which are able to modify their data transfer rates according to available bandwidth within the network is termed *elastic traffic*.

The papers mention are given in chronological order but some general things should be noted in the various schools of research. What we term the Cambridge papers are the works by the authors Kelly, Gibbens, Key, Tan, McAuley, Laevens and Barham. They are interested in the problem of flow control of elastic traffic in the presence of a pricing scheme. The approach they follow is through consideration of two models which crop up in each paper. Low and Lapsley[LL99] consider the same model also. La, Anantharam, Lazar, Orda, Rom, Shimkin consider the routing problem. The models which are examined tend to be very simple networks and results tend to hold under very specific assumptions such as all users of the network having the same source and destination nodes.

Some papers on pricing as regards what future networks may implement are reviewed. They contain great amounts of information on the actual statistics of the Internet as it behaves now and the problems it faced in the past. We restrict ourselves to only a few as the issue discussed is pricing and not any proposed solutions to network models.

Once users are endowed with utility functions to represent their benefits, the discussion of equilibrium points which arise from the models considered comes upon the issue of *fairness*. There is a difficulty in examining fairness without a specific framework. The different types of fairness are important when considering how rates are allocated in an equilibrium. Max-min is the most popular form of fairness which maintains that no users rate can be increased without decreasing the rate of a smaller user. Proportional fairness has been suggested by Kelly as a more appropriate form of fairness. Another form suggested in [MR99], Potential

Delay Minimization, introduces one in which a bandwidth sharing objective (from the view of file transfers, for example) is the minimization of the time delay needed to complete the transfers.

The use of decentralised algorithms has been suggest as a means of flow control for elastic traffic. The advantage in their use is that they are implemented by the end-systems and decide on alterations to current rates based on congestion signals received from resources in the network. The Cambridge papers look at two classes of decentralised algorithms which yield solutions to their model with willingness-to-pay users.

In giving the abstracts, we find it beneficial to add notes to highlight other issues addressed or interesting points. Throughout the papers, we note that the terms *user* and *end-system* have the same meaning and resources in a network are sometimes referred to as *nodes*.

2 The Abstracts

2.1 The Cambridge Papers

Charging and rate control for elastic traffic

F.P. Kelly [Kel97]

Abstract

This addresses the issues of charging, rate control and routing for a communication network carrying elastic traffic, such as an ATM network offering an available bit rate service. A model is described from which max-min fairness of rates emerges as a limiting special case: more generally, the charges the users are prepared to pay influence their allocated rates. In the preferred version of the model, a user chooses the charge per unit time that the user will pay; thereafter the users' rate is determined by the network according to a proportional fairness criterion applied to the rate per unit charge. A system optimum is achieved when users' choices of charges and the networks choice of allocated rates are in equilibrium.

Refers to: [MMV94], [She95]

Notes

As the title describes, a model is presented. The model is of interest as it arises throughout the papers [FKT98], [GK99], [KM99], [PKL99] so it will be useful to describe here.

He considers finite sets of resources \mathcal{J} , routes \mathcal{R} and users \mathcal{S} . A route is a collection of resources and each resource has some finite capacity. They assume that each route is used by exactly one user but a user may operate on several routes. Each user has a utility which is a function of the sum of the loads placed on the routes it operates and the utility functions are assumed to be strictly increasing and convex functions of the rates. A resource charges each of its users in proportion to the loads they contribute, the net charge being an increasing function of the total load upon it.

Kelly looks at the system optimization problem which is to maximize the sum of the users utilities. The problem is examined from two different pricing schemes the resources may implement, namely a price to each user per unit flow and a price per unit share.

Each user has a utility maximization problem which is to maximize the users net utility, i.e. the utility less the cost, where each users is charged a price(λ_s) per unit flow. The revenue optimization problem for the network is to maximize the sum of the users flows in order to maximize the revenue it receives from the users. The main result is summarised in a way which says the unique solutions to the users maximization problems solves the network revenue optimization problem and also solves the system optimization problem.

The fairness of the optimal solutions depends, amongst other things, upon the users utilities and it is shown that logarithmic utility functions give rise to proportionally fair solutions, which favour smaller flows less emphatically.

The alternative situation is considered where the user is charged an amount m_s per unit time and receives a flow proportional to this quantity, rather than by a charge per unit flow. A similar theorem is shown to hold for this situation as does in Theorem 1, which relates the

solution for the user optimization problems to the optimization problems for the system and network.

Cited by: [FKT98], [GK99], [Kel99], [CCK99], [LL99], [MR99].

Rate control in communication networks: shadow prices, proportional fairness and stability

F. P. Kelly, A.K. Maulloo & D.K.H. Tan [FKT98]

Abstract

This paper analyses the stability and fairness of two classes of rate control algorithm for communication networks. The algorithms provide natural generalizations to large-scale networks of simple additive increase/multiplicative decrease schemes, and are shown to be stable about a system optimum characterized by a proportional fairness criterion. Stability is established by showing that, with an appropriate formulation of the overall optimization problem, the network's implicit objective function provides a Lyapunov function for the dynamical system defined by the rate control algorithm. The network's optimization problem may be cast in primal or dual form: this leads naturally to two classes of algorithm, which may be interpreted in terms of either congestion indication feedback signals or explicit rates based on shadow prices. Both classes of algorithm may be generalized to include routing control, and provide natural implementations of proportionally fair pricing.

Keywords: ATM network, congestion indication, elastic traffic, Internet, Lyapunov function, proportionally fair pricing, queues, routing, tatonnement.

Refers to: [RMD91] [MMV94] [Kel97], [She95],

Notes

This paper follows up the material of [Kel97]. The same model is considered along with two rate control algorithms which converge close to the optimal solution of the social optimization problem which was derived in [Kel97]. The specific point to which the algorithms converge is the solution of a relaxation of this model, the authors look at what the problem (in terms of networks) it is the solution to.

The theoretical results are illustrated through examples using a four node network with Poissonian assumptions on the traffic and then a random network with 100 resources and 1000 routes. A very interesting paper from a mathematical viewpoint.

Cited by: [KM99], [GK99], [PKL99], [Kel99], [CCK99], [LL99], [MR99].

Resource pricing and the evolution of congestion control

R.J. Gibbens & F.P. Kelly [GK99]

Abstract

We describe ways in which the transmission control protocol of the Internet may evolve to support heterogeneous applications. We show that by appropriately marking packets at overloaded resources and by charging a fixed small amount for each mark received, end-nodes are provided with the necessary information and the correct incentive to use the network efficiently.

Keywords: charging, critical congestion interval, game theory, Internet, proportionally fair pricing, rate control, sample path shadow prices.

Refers to: [Kel97] [FKT98] [KM99] [MMV94], [CCK99], [CJ89], [Cla96], [Jac88], [She95]

Notes

The authors use a simple slotted model of a resource to illustrate how network shadow prices may be identified. The marking mechanism is for the resource to mark every packet which arrives into an overloaded slot.

They also study how these sample path shadow prices may be used to transfer information and incentives to end-nodes. An investigation into the behaviour of a system with a resource and end-nodes equipped with various transmission control algorithms designed to achieve different user objectives is carried out.

They go on to define sample path shadow prices for a more realistic queueing model of a resource. For a queue though, it is often not possible to be sure whether or not a packet should be marked until some time after the packet has left the resource. They describe several marking mechanisms which attempt to approximate the ideal behaviour.

Finally they compare and contrast the mechanisms of this paper with the current Internet.

Cited by: [KM99], [PKL99], [Kel99], [LL99].

Differential QoS and Pricing in Networks: where flow-control meets game theory

Peter B Key & Derek R McAuley [KM99]

Abstract

This paper looks at ways of providing Quality of Service to users based on a simple pricing scheme. It is primarily aimed at elastic traffic, and it is users rather than the network who define the flow control schemes. A framework for assessing schemes and algorithms via a distributed game is presented.

Refers to: [FKT98], [GK99], [MMV94], [Jac88], [Od199], [She95].

Notes

The starting point of this paper are the papers [FKT98] and [GK99]. The model they introduce is the same as that of [FKT98]. Their idea is to let users generate a load on the network and receive feedback signals related to shadow prices.

If the costs reflect the loss, then packets should be marked if and only if the packet contributed to a loss event. Thus let the network determine the prices and let the users respond and place demands upon the networks which the network will try to meet.

The implementation of these types of flow control algorithms requires some method by which the network may feed information to the users. Many data transfer protocols have a mechanism for congestion indication to which we could fit a form of Explicit Congestion Notification. A suggestion is made of how the TCP algorithm may be modified to incorporate such a mechanism.

Some thought is given to the willingness-to-pay algorithms proposed in [GK99] concerning equilibriums and steady states.

The authors suggest a practical framework in which a distributed emulation of a network (a cluster of PCs) to which users connect across the Internet and run their own algorithms. The details are given at: <http://www.statslab.cam.ac.uk/~richard/research/topics/evolution/>.

They point out that the analysis of these types of control schemes may be viewed from widely diverse disciplines such as: Computer Science, Control Theory, Game Theory, Economics, Financial Mathematics, Stochastic Decision Theory, Evolutionary biology and Optimization / Dynamic Programming. What is considered optimal crucially depends upon what your desires are and some typical objectives which they list might be:

1. To keep the expected rate of cost constant.
2. Maximize the average throughput minus the cost.
3. Maximize throughput given a cost constraint.
4. Maximize the net utility.
5. Transfer a file at minimum cost.
6. Transfer a file within a fixed time at minimum cost.
7. Transfer a file as quickly as possible at minimum cost.

Cited by: [GK99], [PKL99], [Kel99].

Congestion Pricing for Congestion Avoidance

Peter Key, Derek McAuley, Paul Barham & Koenraad Laevens [PKL99]

Abstract

This paper describes the use of Congestion Pricing as a means of providing Congestion Control and Differentiated Quality of Service. The application of the proposed technique to the Internet Protocol has the advantage that it can be simply implemented using Explicit Congestion Notification. In particular: the network mechanism is independent of higher level protocols; the end systems can continue to exhibit current TCP behaviours; new multiprotocol flexibility is made available to end systems and users.

Architectural issues are discussed, including important aspects of aggregation and charging. We describe our methodology for assessing the scheme via a distributed network simulator. Initial results are presented which compare and contrast various adaptive strategies that achieve a variable range of TCP-like behaviours.

Refers to: [GK99], [FKT98], [KM99], [MMV94], [RMD91], [She95].

Notes

The paper builds upon the framework of the rate control algorithm mention in [FKT98] and suggests Explicit Congestion Notification as the form of pricing.

The authors conduct experiments in which they compare the proposed willingness-to-pay rate control algorithm to TCP and TCP with explicit congestion notification. The willingness-to-pay algorithm is seen to be more stable in terms of average throughput than TCP. This is probably due to the small value of the constant κ which controls the effect marks received has on the rate which is being sent. As regards TCP, it reacts to congestion by decreasing its window size by a factor of a half.

Cited by: [Kel99].

Mathematical modelling of the Internet

Frank Kelly [Kel99]

Abstract

Modern communication networks are able to respond to randomly fluctuating demands and failures by adapting rates, by rerouting traffic and by reallocating resources. They are able to do this so well that, in many respects, large-scale networks appear as coherent, almost intelligent, organisms. The design and control of such networks present challenges of a mathematical, engineering and economic nature. This paper outlines how mathematical models are being used to address current issues concerning the stability and fairness of rate control algorithms for the internet.

Refers to: [GK99], [FKT98], [Kel97], [KM99], [PKL99], [LL99], [Jac88], [She95].

Notes

Basically a review paper based on [Kel97],[FKT98] and [LL99]. They provide a discussion on the stability and fairness of a dynamical system modelling TCP and a generalistion MulTCP.

Cited by: None.

An Intelligent Agent for Optimizing QoS-for-Money in Priced ABR Connections

C. Courcoubetis, G.D. Stamoulis, C. Manolakis & F.P. Kelly [CCK99]

Abstract

We investigate the usage and functionality of an Intelligent Agent at the user-side in an environment where multiple users (not necessarily identical) are served by and charged for ABR connections. We formulate an optimization problem pertaining to the case where bandwidth allocation is proportionally fair, and is based on the amount of money users are willing to pay per unit time. Each user is assumed to select his willingness-to-pay so as to maximize his net benefit, namely the difference of the utility acquired by the QoS attained minus the willingness-to-pay. The optimal selection depends on the network state, namely on the total willingness-to-pay by other users and on the total capacity available for ABR connections. We develop an Intelligent agent that replaces the user in choosing the willingness-to-pay, adaptively to the network state, by learning the user preferences. This is based on past history of user choices for a wide range of network states. We analyze the main issues related to the above problem, and we present an algorithm for the Intelligent agent to learn user preferences and select the willingness-to-pay on his behalf. We also provide simulation results showing that the performance of this algorithm is very close to optimal. Our approach can be applied to more general cases of economic sharing of network resources, and offers new capabilities for resource management, while it is not specific to ABR; it can in general be applied to elastic services, such as TCP / IP.

Refers to: [Kel97], [FKT98].

Cited by: [GK99].

Optimization Flow Control, I: Basic Algorithm and Convergence

Steven H. Low & David E. Lapsley [LL99]

Abstract

We propose an optimization approach to flow control where the objective is to maximize the aggregate source utility over their transmission rates. We view network links and sources as processors of a distributed computation system to solve the dual problem using gradient projection algorithm. In this system sources select transmission rates that maximize their own benefits, utility minus bandwidth cost, and network links adjust bandwidth prices to coordinate the sources' decisions. We allow feedback delays to be different, substantial and time varying, and links and sources to update at different times and with different frequencies. We provide asynchronous distributed algorithms and prove their convergence in a static environment. We present measurements obtained from a preliminary prototype to illustrate the convergence of the algorithm in a slowly time-varying environment. **Keywords:** Optimization flow control, congestion pricing, gradient projection, asynchronous algorithm, Convergence

Refers to: [GK99], [Kel97], [FKT98], [She95].

Notes

The authors use the same model as [Kel97]. The same system optimization problem is solved but the authors differ in their approach to the solution which gives rise to different flow control algorithms. This leads to a different marking implementation of the algorithms, essentially users decide on their rates and pay what the network charges, whereas in [Kel97] users decided upon their payments and received whatever rate the network allocated. The approach used was to solve the dual problem using gradient projection method.

Experimental measurements are given from an implementation of the above theory on a network consisting of two PC's connected via ethernet which support the claim that the results hold even when the network conditions vary slowly.

Cited by: [Kel99].

Bandwidth sharing: objectives and algorithms

L. Massoulié & J. Roberts [MR99]

Abstract

This paper concerns the design of distributed algorithms for sharing network bandwidth resources among contending flows. The classical fairness notation is the so-called max-min fairness; F. Kelly [Kel97] has recently introduced the alternative proportional fairness criterion; we introduce a third criterion, which is naturally interpreted in terms of the delays experienced by ongoing transfers. We prove that fixed size window control can achieve fair bandwidth sharing according to any of these criteria, provided scheduling at each link is performed in an appropriate manner. We next consider a distributed random scheme where each traffic source varies its sending rate randomly, based on binary feedback information from the network. We show how to select the source behavior so as to achieve an equilibrium distribution concentrated around the considered fair rate allocations. This stochastic analysis is then used to assess the asymptotic behaviour of deterministic rate adaption procedures.

Refers to: [Kel97], [FKT98], [CJ89], [Jac88].

Notes

This paper sees the introduction of an alternative fairness criterion in which the dynamics of network bandwidth allocations may be examined. As opposed to both max-min and proportional fairness, *potential delay minimization* fairness is sensitive to minimizing time delays in general (for example in the light of file transfers). It is intermediary to the previously mentioned criteria in the sense that it long routes are more severely penalised than max-min fairness, resulting in a larger overall throughput but less hash upon long routes than proportional fairness giving way to a smaller overall throughput.

Explicit rate algorithms generally impose severe constraints on network nodes and rely on uniform implementation throughout the network for optimal efficiency. Thus in view of the complexity of such algorithms, most network flow protocols are based on simple binary indications of congestion issued independently by the network links.

If link buffers are sufficiently large enough to eliminate the possibility of data loss then a feasible bandwidth sharing option is a reliance upon non-adaptive end-to-end windows.

The authors explore how window sizes, round trip times but more importantly the choice of service policy for the links in the network affect the type of fairness which emerges. Thus showing how a simple fixed window control protocol produces different sharings depending upon the scheduling discipline employed at the networks nodes. Fifo is seen to give rise to weighted proportional fairness while Fair Queueing leads naturally to max-min fairness.

Also approached is the problem of designing a distributed algorithm to realize a given sharing objective through the study of a family of so-called Metropolis algorithms.

Cited by: None.

A modest proposal for preventing Internet congestion

Andrew Odlyzko [Odl99]

Abstract

A simple approach, called PMP (Paris Metro Pricing), is suggested for dealing with congestion in packet networks such as the Internet. It is to partition a network into several logical networks, each of which would treat all packets equally on a best effort basis, just as the current Internet does. There would be no formal guarantees of quality of service. The separate networks would differ only in the prices paid for using them. Networks with higher prices would attract less traffic, and thereby provide better service. Price would be the primary tool of traffic management.

Refers to: [Cla96], [MMV94], [She95].

Notes

As the author says in reference to Jonathan Swifts essay, “I propose to turn a perceived burden into a solution, and rely on usage-sensitive pricing to control congestion, by passing most of the complexity of other solutions.” The idea of PMP is interesting in itself and the following description contains the heuristics of the approach:

Until about 15 years ago, when the rules were modified, the Paris Metro operated in a simple fashion, with 1st and 2nd class cars that were identical in number and quality of seats. The only difference was in the price of 1st and 2nd class tickets. The result was that 1st class cars were less congested, since only people who cared about being able to get a seat, not having to put up with noisy teenagers, etc., paid for 1st class. The system was self-regulating, in that whenever 1st class cars became too popular, some people decided they were not worth the extra cost, and traveled 2nd class, reducing congestion in 1st class and restoring the differential in quality of service between 1st and 2nd class cars.

It is thought by certain people in the data networking community that all resources should be devoted to improving capacity (the “fat dumb pipe” model) rather than working on complicated network schemes. The author admits that the PMP proposal is close to the “fat dumb pipe” model but points out that it brings in economic incentives to provide uncongested pipes for those who need them.

The author notes that this approach is only a suggestion and that extensive research would be required before any realistic implementation was in place. The problems which it faces in the current Internet are discussed and thought is put into its implementation. An interesting point is that if it were introduced, at least initially, the cost per packet on the lowest cost network would undoubtedly be zero. A review of other pricing proposals is given at the end.

Cited by: [KM99].

2.2 Routing

Fairness in Network Optimal Flow Control: Optimality of Product Forms

Ravi Mazumdar, Lorne G. Mason & Christos Douligeris [RMD91]

Abstract

In this paper we consider the problem of optimal flow control in a multiclass telecommunications environment where each user (or class) desires to optimize its performance while being fair to the other users (classes). The Nash arbitration scheme from game theory is shown to be a suitable candidate for a fair, optimal operation point in the sense that it satisfies certain axioms of fairness and is Pareto optimal. This strategy can be realized by defining the product of individual user performance objectives as the network optimization criterion. This provides the rationale for considering the product of user powers as has been suggested in the literature. For delay constrained traffic, the constrained optimization problem of maximizing the product of user throughputs subject to the constraints leads to a Nash arbitration point. It is shown that these points are unique in throughput space and we also obtain some convexity properties for power and delays with respect to throughputs in a Jackson network.

Cited by: [FKT98], [PKL99], [AOS93].

A Game Theoretic Perspective to Flow Control in Telecommunication Networks

Christos Douligeris & Ravi Mazumdar [DM92]

Abstract

Multiple classes of traffic with differing and often conflicting requirements arise in an integrated telecommunications environment as users share the limited existing resources. In this paper, a game theoretic perspective is presented and analysed as the appropriate framework for the study of the flow control problem. Using the notion of power as the performance criterion, we compare a network-Pareto optimal solution—with two user optimal solutions—Nash and Stackelberg equilibria. The appropriateness of each solution is discussed given the operating characteristics of the system. A proposed greedy algorithm is shown to converge to the Nash equilibrium.

Refers to: [BL87].

Cited by: [LA98].

Competitive Routing in Multi-User Communication Networks

Ariel Orda, Raphael Rom & Nahum Shimkin [AOS93]

Abstract

We consider a communication network shared by selfish users. Each user seeks to optimize its own performance by controlling the routing of its given flow demand, giving rise to a non-cooperative game. We investigate the Nash equilibrium of such systems. For a two-node multiple-links system, uniqueness of the Nash equilibrium is proved under reasonable convexity conditions. It is shown that this Nash equilibrium point possesses interesting monotonicity properties. For general networks, these convexity conditions are not sufficient for guaranteeing uniqueness, and a counter example is presented. Nonetheless, uniqueness of the Nash equilibrium for general topologies is established under various assumptions.

Refers to: [RMD91], [Laz91], [ES91], [She90].

Notes

This paper considers the basic networking problem or routing in a environment consisting of selfish users. It is modelled as a non-cooperative static game. The routing problem is approached in two phases. Firstly they consider a two node network connected by a set of parallel links. Users choose their routing so as to optimize a certain selfish criterion (such as its own average delay). The existence and uniqueness of a Nash equilibrium is established under fairly weak convexity properties which are satisfied by standard network cost functions. The characteristics of and convergence to such an equilibrium are discussed.

This is followed by a more general network. An example shows that the previously used weak convexity conditions are not sufficient to guarantee the uniqueness of a Nash equilibrium although in several particular cases it is shown that uniqueness can hold. The paper ends with a small discussion on the subject matter and some open problems to be tackled.

Cited by: [LA98],[YAKO97a], [YAKO97b], [LA98].

Achieving Network Optima Using Stackelberg Routing Strategies

Yannis A. Korilis, Aurel A. Lazar & Ariel Orda [YAKO97a]

Abstract

In non-cooperative networks users make control decisions that optimize their individual performance objectives. Nash equilibria characterize the operating points of such networks. Nash equilibria are generically inefficient and exhibit suboptimal network performance. Focusing on routing, a methodology is devised for overcoming this deficiency, through the intervention of the network manager. The manager controls part of the network flow, is aware of the non-cooperative behavior of the users and performs its routing aiming at improving the overall system performance. The existence of *maximally inefficient* strategies for the manager, i.e. strategies that drive the system into the global network optimum, is investigated. A maximally efficient strategy of the manager not only optimizes the overall performance of the network, but also induces an operating point that is efficient with respect to the performance of the individual users (Pareto efficiency). Necessary and sufficient conditions for the existence of a maximally efficient strategy are derived, and it is shown that they are met in many cases of practical interest. The maximally efficient strategy is shown to be unique and it is specified explicitly.

Refers to: [Laz91], [ES91], [AOS93], [YAKO97b].

Cited by: [YAKO97b], [LA98].

Capacity Allocation under Non-cooperative Routing

Yannis A. Korilis, Aurel A. Lazar & Ariel Orda [YAKO97b]

Abstract

The capacity allocation problem in a network that is to be shared by non-cooperative users is considered. Each user decides independently upon its routing strategy, so as to optimize its individual performance objective. The operating points of the network are the Nash equilibria of the underlying routing game. The network designer aims to allocate link capacities, so that the resulting Nash equilibria are efficient, according to some systemwide performance criterion. In general, the solution of such design problems is complex and at times counterintuitive, since adding link capacity might lead to degradation of user performance. For systems of parallel links, we show that such paradoxes do not occur and that the capacity allocation problem has a simple and intuitive optimal solution, that coincides with the solution in the single user case.

Refers to: [Laz91], [ES91], [AOS93], [ALP95], [YAKO97a], [YAKO97b].

Cited by: [YAKO97a], [LA98].

Optimal Routing Control: Game Theoretic Approach

Richard J. La & Venkat Anantharam [LA98]

Abstract

Communication networks shared by selfish users are considered and modeled as noncooperative *repeated* games. Each user is interested only in optimizing its own performance by controlling the routing of its load. We investigate the existence of a Nash equilibrium point (NEP) that achieves the system-wide optimum cost. The existence of a subgame-perfect NEP that not only achieves the system-wide optimum cost but also yields a cost for each user no greater than its stage game NEP cost is shown for two-node multiple link networks. It is shown that more general networks where all users have the same source-destination pair have a subgame-perfect NEP that achieves the minimum total system cost, under a mild technical condition. It is shown general networks with users having multiple source-destination pairs don't necessarily have such an NEP.

Refers to: [DM92], [YAKO97b], [YAKO97a], [AOS93], [BL87], [Laz91], [ALP95], [She90].

Notes

Armed with the same model as [AOS93], but using a dynamic rather than static game theoretic framework, the authors first examine the two node network with a fixed source-destination pair is common to all users. Restricting themselves to what are termed type- C cost functions from [AOS93], they show that there are NEPs where the agents operate at the unique system-wide optimum point, where at the same time, the users cost is no greater than it would be in the unique stage game NEP.

If we consider the more general case where users may have different source-destination pairs, an example shows (as it did in [AOS93]) that the existence of a NEP for the repeated game is not guaranteed. If we consider a class user as a coalition of all users that have a particular source-destination pair, then there always a NEP in the repeated game which achieves a total system cost that is no more than the minimal total system cost over all the NEPs of the static game played between these class users. Another example shows that the system-wide minimum cost could be strictly less than this minimum (in some cases as much as 20%) so that the preceding result is not strong enough.

Cited by: None.

2.3 Pricing

Pricing the Internet

Jeffrey K. MacKie-Mason & Hal R. Varian [MMV94]

Abstract

This paper was prepared for the conference “Public Access to the Internet,” JFK School of Government, May 26–27, 1993. We describe the technology and cost structure of the NSFNET backbone of the Internet, and discuss how one might price Internet access and use. We argue that usage-based pricing is likely to be necessary to control congestion on the Internet and propose a particular implementation of usage-based pricing using a “smart market”.

Refers to: [MMV93]

Notes

As the authors admit, This paper overlaps substantially with the Some Economics of the Internet paper [MMV93]. It has slightly less information on the history and suchlike of [].

Cited by: [Kel97], [FKT98], [GK99], [KM99], [PKL99], [Odl99], [MMV93], [Cla96].

Some Economics of the Internet

Jeffrey K. MacKie-Mason & Hal R. Varian [MMV93]

Abstract

This paper was prepared for the Tenth Michigan Public Utility Conference at Western Michigan University March 25–27, 1993. We describe the history, technology and cost structure of the Internet. We also describe a possible smart-market mechanism for pricing congestion on the Internet.

Keywords: Networks, Internet, NREN, NII.

Refers to: [MMV94].

Notes

As the abstract says, they authors outline the history of the Internet and describe some of the technological and economic issues such as financial costs pertaining to lines and routers routers ,both as they were in the past and as they are now.

Several aspects of a smart market solution are considered such as bid setting, accounting for charges, routing and capacity expansion. A long intermittent discourse on the practical applicability to the current network is outlined in terms of social opinion, financial costs and comparisons of the Internet to various economic models from which their arguments come.

The authors note that several problems would have to be solved before any realistic implementation of the smart market could be executed in a packet-switching network.

Cited by: [MMV94].

Adding service discrimination to the internet

D.D. Clark [Cla96]

Abstract

This paper explores extensions to the Internet that can provide discrimination in the service offered to different users in times of network discrimination. It proposes a scheme which allows different users to adjust their sending rates to different values during overload. This scheme is contrasted with a number of resource allocation schemes under consideration.

Refers to: [MMV94], [Jac88].

Notes

The underlying hypothesis of this paper is the characteristic of Internet service most valued by the user is the overall throughput achieved by a user during the transfer of a data object of some size and not the delay of the individual packets. So the relation between the treatment of individual packets and overall user satisfaction is whether a packet triggers a congestion feedback indication, not how much is delayed. This is because congestion is taken as an indication to the source that it is to slow down its sending rate.

The paper suggests that instead of allocating capacity to users by explicit reservations along a path, we should take the much simpler step of aggregating all the traffic that is within the usage profile of all the users, as indicated by the tags in the packets, and then viewing the successful transport of this aggregated traffic as a provisioning problem.

This paper claims that allowing the user to tag packets, because it represents how resources are allocated when they are in demand, represents a rational basis for cost allocation. The mechanism proposed here has the virtue that it is simple to implement and capable of implementing a wide range of policies for allocation of capacity among users.

Since experience suggests that we will see vary creative pricing strategies to attract users, limiting the knowledge of these to a single point, where the user attaches to the network, is the key to allowing providers to differentiate their services with only local impact.

Thus this approach attempts to minimize what we must agree on and deploy in common throughout the Internet, and leaves as much of the total mechanism as a local matter to each provider.

Cited by: [GK99], [Odl99].

References

- [ALP95] A. Orda A. Lazar and D. Pendarakis. Virtual path bandwidth allocation in multi-user networks. *Proceedings IEEE INFOCOM '95*, 1:312–320, 1995.
- [AOS93] R. Rom A. Orda and N. Shimkin. Competitive routing in multi-user communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, October 1993.
- [BL87] A. Bovopoulos and A.A. Lazar. Decentralised algorithms for optimal flow control. *Proceedings of the Twenty-fifth Annual Allerton Conference on Communication, Control and Computing*, pages 979–988, 1987.
- [CCK99] C. Manolakis C. Courcoubetis, G.D. Stamoulis and F.P. Kelly. An intelligent agent for optimizing qos-for-money in priced abr connections. *Telecommunications Systems, Special Issue on Network Economics.*, 1999.
- [CJ89] D.M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Computer Networks and ISDN Systems*, 17:1–14, 1989.
- [Cla96] D.D. Clark. Adding service discrimination to the internet. *Telecommunications Policy*, 20:169–181, 1996.
- [DM91] Z. Dziong and L. Mason. Fair-efficient call admission control policies for broadband networks – a game theoretic framework. *IEEE/ACM Transactions on Networking*, 4:123–136, December 1991.
- [DM92] Christos Douligeris and Ravi Mazumdar. A game theoretic perspective to flow control in telecommunication networks. *Journal of the Franklin Institute*, 329(2):383–402, March 1992.
- [ES91] A. Economides and J. Silvester. Multi-objective routing in integrated services networks: A game theory approach. *IEEE INFOCOM '91*, 3:1220–1227, 1991.
- [FKT98] A.K. Maulloo F.P. Kelly and D.K.H. Tan. Rate control for communications networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [GK99] R.J. Gibbens and F.P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35, 1999.
- [Jac88] V. Jacobson. Congestion avoidance and control. *Proc. ACM SIGCOMM '88*, pages 314–329, 1988.
- [Kel97] F.P. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- [Kel99] Frank Kelly. *Mathematical modelling of the internet*. 1999.
- [KM99] Peter B. Key and Derek R. McAuley. Differential QoS and pricing in networks: where flow-control meets game theory. *IEE Proceedings Software*, March 1999.

- [LA98] Richard J. La and Venkat Anantharam. Optimal routing control: Game theoretic approach. <http://diva.EECS.Berkeley.EDU:80/~ananth/>, June 1998.
- [Laz91] M. Hsiao & A.A. Lazar. Optimal decentralised flow control of markovian queueing networks with multiple controllers. *Perform. Eval.*, 13:181–204, 1991.
- [LC85] H. Lee and M. Cohen. Multi-agent customer allocation in a stochastic service system. *Manage. Sci.*, 31:752–763, June 1985.
- [LL99] Steven H. Low and David E. Lapsley. Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 1999.
- [MMV93] J. MacKie-Mason and H. Varian. Some economics of the internet. Technical report, 1993.
- [MMV94] J. MacKie-Mason and H. Varian. Pricing the internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice-Hall, New Jersey, 1994.
- [MR99] L. Massoulié and J. Roberts. Bandwidth sharing: objectives and algorithms. *INFOCOM 99*, 1999.
- [Odl99] A. Odlyzko. A modest proposal for preventing internet congestion. <http://www.research.att.com/~amo/doc/>, 1999.
- [PKL99] Paul Barham Peter Key, Derek McAuley and Koenraad Laevens. Congestion pricing for congestion avoidance. Technical Report MSR-TR-99-15, Microsoft Research, 1999.
- [RMD91] Lorne G. Mason Ravi Mazumdar and Christos Douligeris. Fairness in network optimal flow control: Optimality of product forms. *IEEE Transactions on Communications*, pages 775–782, 1991.
- [She90] S. Shenker. Efficient network allocations with selfish users. *Proc. Perform. '90*, pages 279–285, September 1990.
- [She95] S. Shenker. Fundamental design issues for the future internet. *IEEE J. Selected Area Communications*, pages 1176–1188, 1995.
- [YAKO97a] Aurel A. Lazar Yannis A. Korilis and Ariel Orda. Achieving network optima using stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5(1):161–173, February 1997.
- [YAKO97b] Aurel A. Lazar Yannis A. Korilis and Ariel Orda. Capacity allocation under noncooperative routing. *IEEE Transactions on Automatic Control*, 42(3):309–325, March 1997.
- [YS95] K. Yamaoka and Y. Sakai. A packet routing based on game theory. *2nd Asia-Pacific Conference on Communications*, 2:954–958, 1995.
- [YS96] K. Yamaoka and Y. Sakai. A packet routing method based on game theory. *Trnasactions of the Institute of Electronics, Information and Communication Engineers B-I*, J79B-I:73–79, March 1996.