Workshop on Economics of Communication Networks

Invited Talk titles and Abstracts (March 4, 2004)

Edwin K. P. Chong, Colorado State University, echong@engr.colostate.edu

Online Pricing for Bandwidth Provisioning in Multi-class Networks

We consider the problem of pricing for bandwidth provisioning over a single link, where users arrive according to a known stochastic traffic model. The network administrator controls the resource allocation by setting a price at every epoch, and each user's response to the price is governed by a demand function. We formulate this problem as a partially observable Markov decision process (POMDP), and explore two novel pricing schemes---reactive pricing and spot pricing---and compare their performance to appropriately tuned flat pricing. We use a gradient-ascent approach in all the three pricing schemes. We provide methods for computing unbiased estimates of the gradient in an online (incremental) fashion. Our simulation results show that our novel schemes take advantage of the known underlying traffic model and significantly outperform the model-free pricing scheme of flat pricing.

This research is joint work with Uday Savagaonkar and Robert Givan.

Michael Devetsikiotis , North Carolina State University, mdevets@eos.ncsu.edu

Measurement-based Resource Allocation and On-line Scheduling of Differentiated Services

We first present an overview of recent efforts in the area of measurement-based, adaptive resource allocation. Topics range from on-line measurement of "effective bandwidths" and other traffic envelopes, to constrained optimization of resources under given pricing schemes, and to issues of distributed implementation in real networks. We then turn our attention to our current work on adaptive scheduling and we discuss the on-line optimization of scheduler settings as well as the behavior of the solutions in different pricing and utility scenarios.

In collaboration with George Michailidis, Dept. of Statistics, U. of Michigan, and Peng Xu, North Carolina State University.

Ramesh Johari, MIT, rjohari@mit.edu

A game theoretic view of efficiency loss in network resource allocation

In collaboration with S. Mannor and J. Tsitsiklis, MIT.

Scott Jordan , University of California, Irvine, sjordan@uci.edu

Pricing in the Internet

The Internet will provide different qualities of service (QoS) to different applications. Priority service will be offered through Differentiated Services (DiffServ), while guaranteed performance will be offered through Integrated Services (IntServ). Our premise is that it is not sufficient to introduce DiffServ into the Internet without also considering how much traffic is labelled with each priority class. Only by having information about the performance achieved by each priority class can users (or user agents) decide how to label traffic. Similarly, in IntServ there must be a mechanism to for a user to map buffer and bandwidth into QoS and for the network to determine resource allocation among traffic classes.

Our approach is to introduce the notion of pricing. Users (or their agents) and the network negotiate for resource usage by exchanging price and demand information. The prices indicate the optimal benefit gained by bandwidth and buffer usage on each link and route, based on current demand. Prices may either be used purely as internal network indicators of congestion or as real charges to each user. By distributing the resource allocation process using pricing, we aim to combine existing research results on priority service and resource reservations with results on traffic models, QoS metrics, resource management architectures, statistical multiplexing, and connection admission control to provide a complete connection establishment process for the Internet.

George Kesidis, Penn State, kesidis@engr.psu.edu

Pricing residential broadband access for the emerging Internet

Residential penetration of broadband Internet access, both DSL and cable, is growing very rapidly. Furthermore, technology may soon expand available access bandwidth from about a maximum of 10Mbps today to 100Mpbs in the near future. Demand for broadband access is fueled by the desire for "value-added" services (such as interactive gaming, emerging peer-to-peer applications, IP telephony, etc.) and by the affordable additional cost of subscription over dial-up access. We will discuss the commercial and cyber security concerns of such a dramatic

increase in residential broadband access (RBA). A "first hop diffserv" architecture is proposed so that the infrastructure providers of RBA can recover costs from the value-added services they enable, and also address security concerns, by offering a premium service that is more reliable. Specifically, we focus on the packet memories feeding the links connecting the first-PoP layer-3 routers of the RBA provider to the Internet. A differential enqueue policy is devised and the dynamics of user access are studied when there is congestion (excess demand) in the memory.

Work in collaboration with C. Kirjner, McKinsey & Co., New York.

Richard Hyong-Jun La, University of Maryland, hyongla@glue.umd.edu

Global Stability Conditions for Rate Control with Arbitrary Communication Delays

We adopt the optimization framework for the rate allocation problem proposed by Kelly and characterize the stability conditions with arbitrary communication delays. We demonstrate the existence of a fundamental trade-off between users' price elasticity of demand and the responsiveness of resources through a choice of price function. We apply our results to a family of popular utility and resource price functions, and establish exponential stability. We study the effects of non-responsive traffic on system stability and show that the presence of non-responsive traffic enhances the stability of system.

Joint work with Priya Ranjan and Eyad H. Abed (Univ. of Maryland)

Peter Marbach, University of Toronto, marbach@cs.toronto.edu

Bandwidth Provisioning in Contention-Based Access Networks

We study the link congestion problem in contention-based access networks such as wireless LAN's. These networks are characterized by the use of a random access protocol for channel access. We assume adaptive users or applications in the sense that their traffic flows are elastic. We further assume that individuals' service requirements can be characterized by demand functions that are price-sensitive. A price-based rate control strategy, where the price indicates the current congestion state, is considered to control channel congestion.

The effectiveness of price-based rate control is studied in the context of the classic slotted ALOHA model with an infinite population. The price as the new state variable is dynamically updated based on control parameters and the ternary channel feedback. Our results show that under this model stabilization of the ALOHA channel can be achieved. In particular, by using

drift analysis, we prove that the associated Markov chain is positive recurrent. The resulting steady state probability distribution thus characterizes an operating point for the model. Moreover, a desired operating point as such could be selected by proper choice of the control parameters. We also show that service differentiation is realized at the operating point. From a control perspective, where demand functions become predetermined system parameters, such a price-based scheme offers a simple mechanism to provide service differentiation in a best-effort contention-based network.

This is joint work with Clement Yuen.

Vladimir Marbukh, NIST, marbukh@nist.gov

On Price Based Network Provisioning: Challenges and Approaches

It is known that unlimited price competition among service providers for the user demands as well as competition among users for the network resources may result in undesirable equilibrium resource allocation. We are discussing effect of competition among service providers for the user demands on the equilibrium prices and service availability for the users. We are looking at both, Nash and evolutionary concepts of equilibrium. Pricing for social welfare (aggregate utility) optimization has been considered by F. Kelly et. al. We discuss possible extensions of this work to pricing network resources and network contracts for developing distributed network management schemes intended to maximize the aggregate utility.

Joint work with Robert Van Dyck.

Keith Ross, Polytechnic University, ross@poly.edu

The Peer Resource Market

Today many peers (that is, Internet-connected computer systems) possess surplus bandwidth, storage and CPU resources; for example, a peer with a broadband DSL connection typically utilizes only a small fraction of its transmission, storage and computation capacity. When aggregated together across all peers worldwide, these unused resources constitute a huge, untapped resource pool. This resource pool can potentially be used by many applications, including peer-driven content distribution, globally distributed archival storage, massively parallel computation and P2P file sharing. Of course, applications can only harness the resource pool if peers make available their surplus resources to them. In this talk we'll survey our research

on creating an online marketplace where entities - such as peers, companies, users and so on - buy and sell surplus resources.

Jonathan Shapiro, Michigan State University, jshapiro@cse.msu.edu

Incentives for Cooperative Behavior in Peer-to-Peer Systems

Peer-to-peer design has emerged as an important paradigm both for providing network services in the absence of a fixed service infrastructure. Examples include ad hoc networks, application-level multicast, anonymous communication protocols, and numerous peer-to-peer applications. Even in more traditional network settings where a fixed infrastructure exists, network-network interfaces can be viewed as peer-to-peer applications. In an inter-domain routing protocol, for example, each peer might represent a different autonomous system. A distinguishing characteristic of peer-to-peer systems is that each peer simultaneously presents demand to the system and controls a subset of system resources. Although such systems rely on their users to behave cooperatively by allocating resources to each other, natural incentives exist to withhold resources from the system while still consuming service---a behavior known as free-riding. It is doubtful that peer-to-peer systems can function at large scales in the face of widespread free-riding, leading many researchers to explore the possibility of building incentive mechanisms into systems to encourage user cooperation.

This talk will present an overview of recent work (by the presenter and by others) on several classes of incentive mechanisms. One class of mechanisms, known as reputation systems, record peer interactions allowing peers either individually or collaboratively to reward good behavior or isolate free-riders. An alternate class of mechanisms are based on a notion of currency that must be exchanged for service. Since currency is necessary to acquire service, peers are motivated to cooperate in order to obtain the necessary currency. Reputation and currency-based mechanisms are general-purpose in nature and have been adapted to a range of applications including ad hoc networks, anonymous communication and distributed storage. In addition to such generic mechanisms, there are a few examples of application-specific incentive mechanisms.

The study of incentive mechanisms poses many interesting challenges and opportunities for analytical modeling in the search for mechanisms that provide effective incentives with minimal overhead. We will discuss several examples from the recent literature including: how to assess the cooperativeness of other peers, how to incorporate reputation information from untrusted parties, the use of currency-based mechanisms to promote stable group membership, and the impact of anonymity on cooperative behavior.

Ness Shroff, Purdue University, shroff@ecn.purdue.edu

Simplifications of Network Dynamics in Large Systems

In the first part of the talk, we will show how that significant simplicity can be achieved for pricing-based control of large networks. We consider a general loss network with Poisson arrivals and arbitrary holding time distributions. In dynamic pricing schemes, the network provider can charge different prices to the user according to the current utilization level of the network and also other factors. We show that, when the system becomes large, the performance (in terms of expected revenue) of an appropriately chosen static pricing scheme, whose price is independent of the current network utilization, will approach that of the optimal dynamic pricing scheme. Further, we show that under certain conditions, this static price is independent of the route that the flows take. We then extend the result to the case of dynamic routing, and show that the performance of an appropriately chosen static pricing scheme with bifurcation probability determined by average parameters can also approach that of the optimal dynamic routing scheme when the system is large. These results deepen our understanding of pricing-based network control. In particular, they provide us with the insight that an appropriate pricing strategy based on the average network conditions (hence, slowly changing) could approach optimality when the system is large.

In the second part of the talk, we will also describe how discuss how similar large-system results can be exploited to develop a purely distributed optimization approach for Quality of Service routing in high-bandwidth networks. We develop a distributed and adaptive algorithm that can efficiently solve the optimization online. Compared with existing QoS routing schemes, the proposed optimization based approach has the following advantages: (1) The computation and communication overhead can be greatly reduced without sacrificing performance; (2) The operating characteristics of the network can be analytically studied; and (3) The desired operating point can be tuned by choosing appropriate utility functions.

Jun Shu, School of Business, Penn State, junshu@psu.edu

Creating Label Switched Paths through an Iterative Combinatorial Double-Sided Auction Mechanism

We propose an iterative combinatorial double auction for trading resources in LSP (Label Switched Path) networks. Our goal is to let the market determine an efficient allocation of the resources, and decide a price—under various equilibrium concept—for each link of the network. In a multi-domain network, an end-to-end connection may traverse links that are operated by different owners. Such is the case of the Internet backbone. The current practice is for these owners to conduct secret bilateral negotiations. The resulting bilateral interconnection agreements produce the overall connectivity of the entire multi-domain network. Does there exist a more efficient approach then the secret bilateral negotiation? We think the answer is positive and set

out to design an auction-based market approach to test our hypothesis. The items for auction are resource capacities over the links of the network. This paper treats bandwidth as the only resource. For different end-to-end connection paths, these items may be complementary or supplementary. Traders are owners of these links who may buy and sell capacities at the same time. The end result of the auction is twofold: an efficient allocation of the bandwidth to different traders, and an equilibrium price for each link. The auction is multi-round so as to allow for a more graceful price discovery process.

In collaboration with P. Varaiya, U.C. Berkeley

Rayadurgam Srikant, University of Illinois at Urbana-Champaign, srikant@ifp.uiuc.edu

Economics of Network Pricing with Multiple ISPs

In this talk, we will examine how transit, access and end-user prices are set in a network consisting of multiple ISPs. Some ISPs may be geographically co-located so that they compete for the same set of end users while other ISPs do not directly compete for users, but are nevertheless involved in a noncooperative game of setting access and transit prices for each other. We examine the existence of equilibrium price strategies in these situations and show that positive profit can be achieved using threat strategies.

Joint work with Srinivas Shakkottai.

Kevin Tang, Caltech

A Model of Heterogeneous TCP Algorithms

TCP/AQM protocols can be interpreted as distributed primal-dual algorithms over the Internet to maximize aggregate utility. This interpretation, however, is valid only when all TCP sources react to the same type of price signal (congestion measure) at the links in their paths. We extend the current model to the case where heterogeneous TCP sources that react to different price signals share the same network. This model is useful in understanding the interaction of loss-based TCP algorithms such as Reno and delay-based TCP algorithm such as FAST.

Work in collaboration with Steven H. Low, Caltech, slow@caltech.edu.

Jean Walrand, UC Berkeley, wlr@eecs.berkeley.edu

Issues in Pricing Internet Services

One of the critical challenges facing the telecommunications industry today is to increase the profitability for Internet service providers. For historical reasons, the current Internet protocol stack lacks basic features needed to implement efficient economic mechanisms. Consequently, the providers have limited economic incentives to invest in new technology for value-added services. This results in a stagnant industry and limits the evolution of the Internet.

In this talk, we present pricing schemes that would enable the providers to profit from offering differentiated services and share the increased revenues fairly. We first show that with the commonly accepted Differentiated Services model, if prices are not properly differentiated with respect to service quality, then the system may settle into either unstable or inefficient equilibria. We then discuss how to construct pricing schemes that are stable and lead to socially optimal allocation among users.

Pricing issues become more complex when a service needs to be jointly provided by a network of providers. We first show that if providers are allowed to charge freely in their own interest, then the resulting equilibrium could be inefficient, unfair and may discourage future upgrades to the networks. As an alternative, a simple revenue-sharing policy, under which providers would agree to collaborate for increasing their revenues, can be shown to eliminates all the aforementioned drawbacks. For its implementation, a protocol is constructed based on the predicted outcome of the game, so that providers do not have incentive to cheat. We construct a decentralized algorithm that the providers can use to compute the optimal prices and show that the algorithm is scalable and converges to the improved equilibrium.

In collaboration with Linhai He.
