

# A Survey of Smart Data Pricing: Past Proposals, Current Plans, and Future Trends

SOUMYA SEN, CARLEE JOE-WONG, SANGTAE HA, and MUNG CHIANG,  
Princeton University

Traditionally, network operators have used simple flat-rate broadband data plans for both wired and wireless network access. But today, with the popularity of mobile devices and exponential growth of apps, videos, and clouds, service providers are gradually moving toward more sophisticated pricing schemes. This decade will therefore likely witness a major change in the ways in which network resources are managed, and the role of economics in allocating these resources. This survey reviews some of the well-known past broadband pricing proposals (both static and dynamic), including their current realizations in various consumer data plans around the world, and discusses several research problems and open questions. By exploring the benefits and challenges of pricing data, this article attempts to facilitate both the industrial and the academic communities' efforts in understanding the existing literature, recognizing new trends, and shaping an appropriate and timely research agenda.

Categories and Subject Descriptors: C.2.0 [General]: Data Communications—*Internet, Pricing*; C.2.3 [Network Operations]: Network Management—*Congestion Management*; K.6.0 [Management of Computing and Information Systems]: General—*Economics*

General Terms: Broadband Networks, Network Economics, Congestion Management

Additional Key Words and Phrases: Congestion, broadband pricing, data plans, economics

## ACM Reference Format:

Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. 2013. A survey of smart data pricing: Past proposals, current plans, and future trends. *ACM Comput. Surv.* 46, 2, Article 15 (November 2013), 37 pages.  
DOI: <http://dx.doi.org/10.1145/2543581.2543582>

## 1. INTRODUCTION

In November 1996, AOL, the largest U.S. Internet Service Provider<sup>1</sup> (ISP) of the time, switched from hourly pricing to a flat-rate monthly data plan of \$19.95 [Lewis 1996]. One week later, Pacific Telesis Enterprises CEO M. Fitzpatrick made an apocalyptic prediction [Marshall 1996]:

*My indicators point to a genuine data tsunami, a tidal wave, in the next 18 months. And, ladies and gentlemen, while we can surf the net, nobody can surf a tidal wave.*

<sup>1</sup>An *Internet Service Provider* is defined as a company offering its customers subscriptions for wired and/or wireless access to the Internet (e.g., AT&T, Comcast, and Verizon in the United States).

This work is in part supported by the National Science Foundation, under grant NSF CNS-1117126, and C. J.-W.'s NDSEG fellowship.

Authors' addresses: S. Sen, Carlson School of Management, University of Minnesota; S. Ha, Department of Computer Science and Interdisciplinary Telecom Program, University of Colorado; M. Chiang, Electrical Engineering Department, Princeton University; C. Joe-Wong, Program in Applied & Computational Mathematics, Princeton University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

© 2013 ACM 0360-0300/2013/11-ART15 \$15.00

DOI: <http://dx.doi.org/10.1145/2543581.2543582>

Fortunately, that devastating tsunami did not come, but the tides of growth in the demand for data did help the Internet grow.

However, today, the fear of a data tsunami is slowly coming back [Carew 2012]. The Cisco Visual Networking Index [Cisco Systems 2012a] predicted in May 2012 that global Internet Protocol (IP) traffic will reach 1.3 zettabytes/year (1 zettabyte =  $10^{21}$  bytes) and that mobile data will reach 798 exabytes/year (1 exabyte =  $10^{18}$  bytes) by the end of 2016. With the growing popularity of iPhones, iPads, bandwidth-hungry applications, and cloud-based services, ISPs have been increasingly turning to pricing as the ultimate congestion management tool to mitigate this growth—a feat often achieved by imposing harsh overage penalties and “throttling” the very customers who drive this demand.

The basic idea of congestion-based pricing has existed for several decades in many different utility markets, such as transportation and electricity distribution networks. As for the Internet, even as early as 1997, when it was just evolving into a commercial service and its revenue models were hotly debated, MacKie-Mason et al. [1997] wrote:

*We argue that a feedback signal in the form of a variable price for network service is a workable tool to aid network operators in controlling Internet traffic. We suggest that these prices should vary dynamically based on the current utilization of network resources.*

They were proposing a form of dynamic congestion pricing for the Internet. As a mathematically oriented research topic, congestion pricing has been extensively studied in transportation networks, energy markets, telephone networks, ATM and IP networks, and so forth. Thus, in some sense, there are not many new statements to be made about the generic theory of congestion pricing. And yet, wireless and wireline data networks have so far used the most rudimentary forms of congestion pricing, such as usage-based charges. This can partly be explained by Clark’s [1997] observation:

*Whatever the providers may prefer to do, competition may force some forms of fixed-fee pricing in the marketplace.*

Whereas providers may have preferred flat rates in the 1990s, the situation is quite different now: the problem has worsened and operators are more aggressive in pursuing new pricing schemes. The acceptance of this fact is perhaps nowhere more clearly evident than in the “New Rules for an Open Internet” [Genachowski 2010] announced on December 21, 2010, by Federal Communications Commission (FCC) chairman J. Genachowski:

*The rules also recognize that broadband providers need meaningful flexibility to manage their networks to deal with congestion. . . . And we recognize the importance and value of business-model experimentation.*

Therefore, the interesting question is: given the rapid rise in capacity demand in the age of apps and clouds, how will pricing policies change over this decade?

The shift in pricing trends is more easily noticeable in growing economies, such as India and Africa, where dynamic congestion pricing for voice calls is already practiced. The next logical step in this pricing evolution is dynamic pricing for data. But pricing data has several unique challenges (and opportunities), arising from both technological and social considerations. A clear understanding of these issues requires consideration of both the past and the present developments. To this end, this article reviews some of the best-known pricing proposals of the past two decades and reports on some very interesting pricing schemes that are in use today, along with the identification of research challenges and emerging trends in access pricing.

### 1.1. Contributions

In 1996, Breker [1996] conducted a survey of computer network pricing schemes that covered a few of the known protocols of the time. Similar reviews of pricing concepts for broadband IP and integrated services networks were undertaken about a decade ago by Falkner et al. [2000] and Chang and Petr [2001]. Although these surveys mentioned many concerns that exist even today, such as the need to manage network congestion, since then, researchers have proposed several new schemes for pricing data, and some variations of these plans have been adopted by wireless ISPs for their consumers. Ezziane [2005] surveyed the charging and billing mechanisms used by 3G wireless service providers. In particular, Ezziane reported on new approaches needed for mediation, billing, and charging for these services. More recently, Gizelis and Vergados [2011] provided a detailed overview of the literature on pricing schemes for wireless networks, whereas Sen et al. [2012] studied pricing approaches that incentivize time-shifting of demand. An annotated bibliography of several papers related to various aspects of Internet economics, including pricing, was compiled by Klopfenstein [2009]. Our work not only complements these earlier surveys and reports on broadband pricing but also attempts to understand the kinds of pricing plans currently offered by wired and wireless ISPs around the world while noting some of the weaknesses in their approaches. Additionally, we provide a detailed discussion on the various technological, socioeconomic, and regulatory challenges to pricing innovation, and identify several directions for future research. Our focus in this article is on the pricing for the end user/consumer rather than the economics of peering or transit-fee agreements between service providers. The key features of this work are summarized next:

- Identification of the threats to the Internet ecosystem from the perspective of ISPs, consumers, and content providers.
- Detailed analysis of the various challenges to access pricing innovation and open problems for future research.
- A broad overview of several past and recent proposals for pricing data, along with a classification of static and dynamic pricing plans.
- Several illustrative examples of real pricing plans implemented by ISPs in different parts of the world.
- An overview of the relationships between the ideas of broadband data pricing and earlier congestion pricing ideas in road and electricity networks.
- Discussion on the other directions in access pricing innovation, including satellite and heterogeneous networks.

Networking researchers and operators, caught in their conscientious efforts to find the best solutions to this growing problem, often tend to overlook many of the innovative practices already deployed by network operators in the real world. However, appraising the past several years of theoretical foundations on pricing and identifying their varied realizations in the present world will likely be key to predicting future trends and shaping an appropriate research agenda. The information presented in this work draws from a wide range of sources, including research proposals, existing data plans, news articles, consumer forums, and reports from experimental field trials.

This article is organized as follows: Section 2 provides an overview of threats to the Internet ecosystem from network congestion. An overview of the role of pricing in achieving various network functionality goals, the fundamentals of network economics, and the potential benefits of access pricing schemes in alleviating network congestion is provided in Section 3. Section 4 identifies the key challenges to overcome in order to innovate access pricing and outlines directions for future research. A review of the proposals and realizations of several static and dynamic pricing plans is provided in

Sections 5 and 6, respectively. The emerging trends in network pricing evolution and their roles in satellite and heterogeneous networks are addressed in Section 7. Section 8 concludes with a summary of this work.

## 2. THREATS TO THE INTERNET ECOSYSTEM

The Internet ecosystem today suffers from threats to both its *sustainability* and *economic viability*. The sustainability concern is due to the near doubling of demand for mobile data every year [Cisco Systems 2012a], which brings into question the feasibility of finding any purely technical solutions to this problem. Additionally, changes in the policy to make more spectrum available, although useful, is unlikely to provide any long-term solution. For example, the U.S. government's initiative in releasing an additional 50MHz of spectrum into the market for which six carriers may compete is widely viewed as being quite insufficient [Sachson 2011].

The latter concern regarding the economic viability of Internet access is due to the large overage charges and throttling that ISPs (e.g., AT&T and Verizon) are now imposing on their customers to penalize demand. To understand the implications of these factors, we need to first look at the impact that this growth is having on the various stakeholders of the Internet ecosystem before delving deeper into the promises that Smart Data Pricing (SDP) holds [SDP 2012]. In the following discussion, we explore the current challenges and threats from the perspective of carriers, consumers, and content providers.

### 2.1. ISPs' Traffic Growth

The Cisco Visual Networking Index [Cisco Systems 2012a] predicted in 2012 that mobile data traffic<sup>2</sup> will grow at a Compound Annual Growth Rate (CAGR) of 78% between 2011 and 2016, reaching a volume of 10.8 exabytes per month in 2016. By that time, Cisco predicted that the average smartphone user will consume 2.6GB per month, as opposed to 150MB/month in 2011 [Cisco Systems 2012b]. This growth in per-device data consumption is fueled by the demand for bandwidth-intensive apps: for example, Allot Communications reports that usage of Skype's video calling service grew by 87% in 2010, whereas usage of Facebook's mobile app grew by 267% [Allot MobileTrends 2010]. Devices themselves are also contributing to the increase in traffic volume; Apple's iPad 3 quadrupled the screen resolution from the iPad 2, allowing videos of higher quality to be streamed to the device [Apple 2012].

Cisco predicted that by 2016, WiFi and mobile data will comprise 61% of all Internet traffic, with wired traffic comprising the remaining 39% [Cisco Systems 2012a]. Although the faster growth rate in mobile data traffic is a major concern, growing congestion on wired networks both at the edge and in the networks' middle mile also poses a significant problem. By 2016, ISPs are expected to carry 18.1 petabytes per month in managed IP traffic.<sup>3</sup> That this growth is causing concern among ISPs can be seen from Comcast's recent initiative to cap their wired network users to 300GB per month [Comcast 2012]. Indeed, even in 2008, Comcast made headlines with their (since reversed) decision to throttle Netflix as a way to curb network congestion [Key 2010]. Video streaming due to services like Netflix is in fact a major contributor to wired network traffic—by 2016, Cisco predicts that fixed IPs will generate 40.5 petabytes of Internet video per month [Cisco Systems 2012a].

---

<sup>2</sup>“Mobile data traffic” includes handset-based data traffic, such as text messaging, multimedia messaging, and handset video services, as well as traffic generated by notebook cards and mobile broadband gateways.

<sup>3</sup>Cisco's definition of “managed IP” includes traffic from both corporate IP wide area networks and IP transport of television and video-on-demand.

Congestion in wired networks also affects Rural Local Exchange Carriers (RLECs), due to the persistence of the middle mile problem for RLECs. Although the cost of middle mile bandwidth has declined over the years due to an increase in the DSL demand needed to fill the middle mile, the bandwidth requirements of home users have increased sharply [Glass et al. 2012]. The FCC has set a broadband target rate of 4Mbps downstream speed for home users, but the average speed provided to rural customers is much lower. The cost of middle mile upgrades to meet the target speeds will be substantial and could be a barrier to providing greater speeds for subscribers and digital expansion in the rural areas [Glass et al. 2012]. Therefore, research on access pricing to bring down middle mile investment costs by reducing the peak capacity and RLECs' overprovisioning needs will be a crucial step in bridging the digital divide.

## 2.2. Consumers' Cost Increase

With the growth in the volume of Internet traffic, ISPs have begun to pass some of their network costs on to consumers. Consumers, on the other hand, have become increasingly concerned at the high costs of Internet subscriptions. For instance, when AT&T and Verizon announced in July 2012 that they were offering shared data plans for all new consumers and discontinuing their old plans, many consumers ended up with higher monthly bills [Chen 2012]. Users, in fact, are aware of these increasing costs; while smartphone adoption is still growing in the United States, so is the popularity of usage-tracking and data compression apps that help mobile users stay within their monthly data caps and avoid overage fees (e.g., Onavo, WatchDogPro, DataWiz) [Princeton EDGE Lab 2012]. And such trends are not exclusive to the United States; in South Africa, for instance, consumers' pricing plans for wired networks generally include monthly data caps. To avoid overage charges, many use ISP-provided usage-tracking tools [Chetty et al. 2012]. Even in the United States, research on in-home Internet usage has shown that many users are concerned about their wired Internet bills and would welcome applications for tracking their data usage and even controlling bandwidth rates on in-home wired networks [Chetty et al. 2011]. In fact, although they are not as widely known as similar mobile applications, some apps are offered for wired platforms that allow users to do exactly that [Locktime Software 2012].

## 2.3. Content Providers' Worries

As large U.S. ISPs like AT&T, Comcast, and Verizon move to eliminate flat-rate plans in favor of tiered data plans with usage-based overage fees in both wired and wireless networks, these changes have triggered renewed debate on the Internet's net-neutrality and openness. These discussions have centered around who should pay the price of congestion (i.e., content providers or consumers) and how such pricing schemes should be implemented (i.e., time-of-day, app-based bundles). The major concerns in enabling such pricing practices arise from the possibility of paid prioritization of certain content providers' traffic, price discrimination across consumers, and promoting anti-competitive behavior in bundled offerings of access plus content.

Although such developments can indeed hurt the Internet ecosystem, one aspect that should receive more attention from researchers is the threat to data usage even under current data plans. With Internet users becoming more cautious about the increase in their monthly bills [Teitell 2012], content providers are providing new options to downgrade the Quality of Experience (QoE) for their users to help them save money. For example, Netflix has started allowing "users to dial down the quality of streaming videos to avoid hitting bandwidth caps" [Newman 2011]. Additionally, it is "giving its iPhone customers the option of turning off cellular access to Netflix completely and instead rely on old-fashioned Wi-Fi to deliver their movies and TV shows"



[Fitchard 2012]. Thus, the ecosystem is being driven by an attitude of penalizing demand and lessening consumption through content quality degradation.

Researchers have started to investigate these issues broadly along two lines of work: opportunistic content caching, forwarding, and scheduling; and budget-aware online video adaptation. Opportunistic content delivery involves the smart utilization of unused resources to deliver higher QoE; for example, to alleviate the high cost of bulk data transfers, Marcon et al. [2012] proposed utilizing excess bandwidth (e.g., at times of low network traffic) to transmit low-priority data. Since this data transmission does not require additional investment from ISPs, they can offer this service at a discount, relieving data transfer costs for clients. Although utilizing excess bandwidth introduces some technical issues (e.g., the potential for resource fluctuations), a prototype implementation has shown that they are not insurmountable [Laoutaris et al. 2011]. Hosanagar et al. [2005] examined the pricing of caching in Content Delivery Networks (CDNs), finding that a profit-maximizing operator should offer two vertical classes of service—premium, for which there is some charge, and a free best-effort traffic class. Economic models have also been developed for altruistic cloud providers—that is, those whose main concern is to deliver high-quality service at low cost in response to user queries, instead of maximizing profit [Dash et al. 2009]. Simulations of user queries that mimic those for a scientific database showed that users are able to trade off cost for faster response times.

In a similar spirit, recent works on online video adaptation systems, such as Quota Aware Video Adaptation (QAVA), have focused on sustaining a user's QoE over time by predicting her usage behavior and leveraging the compressibility of videos to keep the user within the available data quota or her monthly budget [Chen et al. 2012]. The basic idea here is that the video quality can be degraded by nonnoticeable amounts from the beginning of a billing cycle based on the user's predicted usage so as to avoid a sudden drop in QoE due to throttling or overage penalties when the monthly quota is exceeded. A common feature of these new research directions is that they address the concerns of consumers and content providers by accounting for both technological and economic factors.

### 3. PRICING DATA NETWORKS

In this section, we first discuss the role that access pricing in data networks can play in achieving different functionality goals and in balancing trade-offs between conflicting goals. We then provide a review of some fundamental concepts and terminology used in the network economics literature to help readers understand the discussions on different pricing proposals in the later sections of this article. This is followed by an overview of the important challenges and open questions in pricing research.

#### 3.1. Network Features

Data networks provide connectivity among users and ubiquitous access to content, which require them to meet a range of functionality goals, such as efficiency, fairness, reliability, manageability, complexity, and implementability. But these goals can often conflict to various degrees. For example, the trade-off between fairness and efficiency in allocating bandwidth to users sharing distinct subsets of links in a data network has been addressed in recent works on multiresource allocation [Joe-Wong et al. 2012b]. Pricing as a mechanism to enable a proportionally fair allocation of bandwidth among network users has also received some attention [Yaïche et al. 2000; Kelly et al. 1998], but more research is needed to understand how pricing can be used a mechanism to achieve a desirable trade-off between fairness and efficiency. Similarly, efficiency, fairness, and reliability goals can be hard to achieve simultaneously, although pricing-based solutions can play a role in balancing these requirements. For instance, Wagner

et al. [2012] developed a Nash Bargaining-based congestion pricing mechanism in a military network setting to distribute jobs across multiple computing clusters for fair load balancing while using redundancy in job replication to provide reliability. The goals of manageability, complexity, and implementability can also conflict; measures to make network traffic manageable usually introduce some system complexity and implementation challenges. Auction-based resource allocation, such as a Smart Market [MacKie-Mason and Varian 1995], and congestion pricing for bandwidth allocation [Murphy and Murphy 1994; Ha et al. 2012a] both improve efficiency and manageability but require additional client-side implementation and increase billing complexity for the service provider. Pricing may also be used as a form of power control; Saraydar et al. [2002] proposed a game-theoretic framework in which wireless users compete for power so as to maximize their Quality of Service (QoS). In summary, various forms of pricing have been proposed by researchers to attain some level of a trade-off between one or more conflicting goals in data networks.

The role that pricing can play in networked systems is not just limited to the Internet. Other utilities like electricity and transportation networks have used pricing mechanisms over the years to regulate demand, such as varying off-peak and peak hour prices or tolls for network access. But access pricing until recently remained largely a topic of theoretical interest in the context of the Internet, and ISPs are only now actively exploring pricing options to manage demand. In fact, several features of data networks differentiate its pricing paradigm from that of other networks. First, unlike traditional electricity grids that involve power generation and its delivery in a point-to-multipoint manner, the Internet is much more decentralized and allows for a bidirectional flow of traffic (i.e., uploads and downloads). As applications and services that require bulk data transfer from the client to the cloud (e.g., machine-to-machine communications, file backup) gain popularity, pricing can be used as a signal for regulating not only the demand for content but also the inflow of content from the edge to the core of the network. Second, the ease of two-way communication between an ISP and its customers (due to well-designed user interfaces for most network-enabled devices) creates the potential for pricing innovations that involve real-time user reaction to feedback signals from the network. In contrast, most household devices connected to the electricity network do not have usage-monitoring capabilities and have to be controlled and scheduled by a centralized controller, as in smart grids [Joe-Wong et al. 2012b]. Third, pricing can exploit the time elasticity of demand more efficiently in the case of Internet services, as many online applications (e.g., movie downloads, peer-to-peer traffic, software updates) are elastic and can be completed without user intervention in small chunks whenever the prices are low [Ha et al. 2012b]. Fourth, consumer billing, privacy, and security issues associated with Internet access pricing, especially for mobile users, are much more complex and sensitive than for electricity and transportation networks. Thus, while pricing can be useful in improving the fairness, efficiency, and manageability of resources, networking researchers also need to account for the additional complexity introduced by the considerations mentioned previously.

### 3.2. Network Economics Fundamentals

In this section, we introduce some of the microeconomic theory concepts [Nicholson and Snyder 2008] commonly used in the broadband pricing literature. We provide a simple example to illustrate several key ideas and terminology. Consider a last-mile ISP that wishes to build a single access link serving a fixed customer base of  $N$  users. The ISP will charge each user based on the amount of bandwidth reserved by that user on this shared access link. Thus, the ISP wishes to determine (1) the optimal capacity  $C$  needed for the access link and (2) the usage-based price  $p$  that should be charged for each unit of bandwidth reserved on this link.




Fig. 1. Computation and interaction of user demand and ISP supply functions in the simplified marketplace of a single access link.

In the traditional parlance of microeconomic theory, the ISP is a *seller* and the users are *buyers* of a particular *good* (i.e., bandwidth). We first consider the individual user's perspective in illustrating the economic model. We index the users by the variable  $i$ ,  $i = 1, 2, \dots, N$ , and denote the bandwidth reserved by user  $i$  as  $x_i$ . Each user is assumed to have a *utility function*  $U_i(x_i, p)$  that represents the satisfaction that a user receives from  $x_i$  amount of bandwidth, after paying a rate  $p$  for unit bandwidth. Given a price  $p$  per unit bandwidth, the user can choose the amount of bandwidth used ( $x_i$ ) so as to maximize  $U_i(x_i, p)$ ; we then obtain a function  $x_i^*(p)$ , which gives the user's optimal bandwidth usage as a function of the price. This function  $x_i^*(p)$  is then called the *user's demand function*. We sum the demand functions across all users to define the aggregate demand function,  $D(p) = \sum_{i=1}^N x_i^*(p)$ , which is the total bandwidth demand as a function of the price. This utility optimization and a schematic demand function  $D(p)$  are illustrated in the left side of Figure 1.

Next, we consider the ISP's perspective. We use  $c(x)$  to denote the cost of building a link with total bandwidth capacity  $x$ . Given a price  $p$  and assuming full utilization of the link, the ISP's utility or *profit function* is then equal to  $\Pi(x, p) = px - c(x)$ , or revenue  $px$  less the incurred cost. The ISP can then choose the capacity  $x^*$  of the network so as to maximize its profit. Since this optimal capacity depends on the price, we denote this function as  $S(p)$ , or the *supply-side function*. The ISP's profit optimization is illustrated on the right side of the diagram in Figure 1. The individual users and the ISP then interact in the marketplace to determine the capacity  $C$  of the single-link network and the price  $p$ . In the center of Figure 1, this interaction is represented by showing the supply function  $S(p)$  and demand function  $D(p)$  on the same axes. The *equilibrium* operating point then occurs when supply equals demand—that is,  $S(p)$  and  $D(p)$  intersect. At this price  $p^*$ , each user will demand and receive her utility-maximizing bandwidth usage, and the total bandwidth  $x^* = C$  available in the network will be fully utilized. The *social welfare*,  $W$ , from this process is calculated by summing up over the utilities of all users and the ISP, or  $W = \sum_i U_i(x_i^*, p^*) + \Pi(x_i^*, p^*)$ .



### 3.3. Pricing as a Congestion Management Tool

The essential goal of smart data pricing is to incentivize users to adjust their behavior so as to better allocate ISP resources among the users. Better resource management in turn leads to a decrease in ISP cost and an increase in profit. Yet using pricing as a congestion management tool requires ISPs to think carefully about users' responses to the prices offered. In particular, the question of timescale often arises when designing pricing plans—should the prices change as network traffic volume changes? And if so, how often and by what amount?

*Static pricing* changes prices on a relatively long timescale (e.g., months or years)—that is, the offered prices do not vary with the network congestion level. The popularity of these plans arises from the certainty they provide regarding users' monthly bills. For instance, tiered data plans with prespecified rates are prevalent in the United States [Frakes 2010; Segall 2011], whereas usage-based pricing in which users are charged in proportion to their usage volume is practiced by several European and Asian ISPs [Parker 2010; Morgan 2011]. But even usage-based pricing leaves a timescale mismatch: ISP revenue is based on monthly usage, but peak-hour congestion dominates its cost structure (e.g., network provisioning costs increase with the peak-hour traffic). Another well-known pricing plan is Time-of-Day (ToD) pricing, in which users are charged higher prices for usage during certain “peak” hours of the day, in order to relieve congestion at these times [Parris et al. 1992]. However, even with ToD pricing, the hours deemed as peak are fixed, which results in two challenges. First, traffic peaks arise in different parts of the networks today, which may end up creating two peaks during the day—one during peak periods for traffic that cannot wait for several hours for lower-price periods, and another peak during discounted off-peak periods for time-insensitive traffic [The Economist 2009]. We discuss several of these existing static pricing plans and proposals in greater detail in Section 5.

*Dynamic pricing* goes further than ToD pricing and does not preclassify peak and off-peak periods, instead adjusting prices at a finer timescale in response to the network congestion level. For instance, in MacKie-Mason and Varian's auction-like Smart Market approach, users “bid” on packets, with the bids reflecting their willingness to pay to send their packets through the network at the given time; the gateway then admits packets with the highest bids so as to maintain network efficiency and charges users according to the minimum successful bid [MacKie-Mason and Varian 1995]. Congestion-dependent pricing adapts the price charged in response to the observed congestion on the network, with higher prices corresponding to higher congestion levels [Kelly et al. 1998; Paschalidis and Tsitsikilis 1998]. However, prices that fluctuate depending on the current network load may be inconvenient for users. Hence, another variant on dynamic ToD pricing, known as the day-ahead pricing, has been proposed to guarantee the prices a day in advance to give users some certainty about the future prices on offer. Each day, new prices are computed for different times (e.g., hours) of the next day, based on predicted congestion levels [Ha et al. 2012b; Joe-Wong et al. 2011]. A detailed discussion on these dynamic pricing proposals will be provided in Section 6.

The previous discussion raises several fundamental questions on the timescale of broadband data pricing. Should prices be based on real-time congestion? Should these price computations be done online or offline to ensure system scalability? Whereas real-time and other forms of dynamic pricing allow ISPs to quickly react to congestion and manage their network resources, rapidly fluctuating prices may not be convenient for consumers, requiring an automated system for submitting bids or reacting to price changes on the client side and a scalable monitoring and price computation system on the ISP side. These and other research challenges are discussed in the next section.

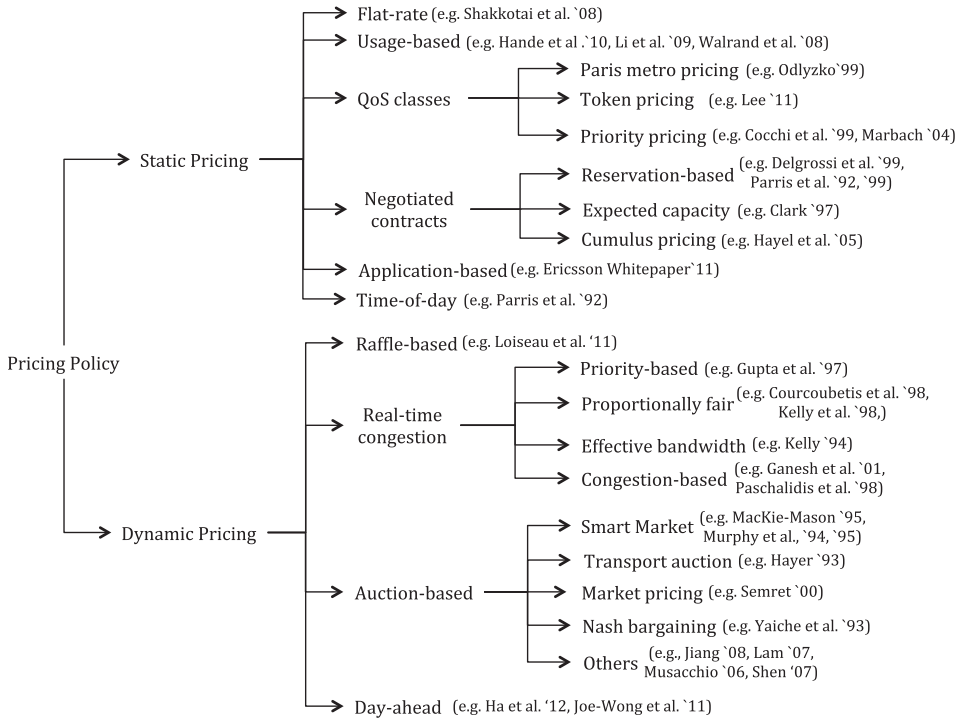


Fig. 2. Taxonomy of selected pricing strategies and some related proposals.

Figure 2 summarizes the pricing plans covered in this survey article, including the static and dynamic plans mentioned earlier. Before discussing the details of these plans, however, we turn to the research challenges posed by broadband data pricing.

#### 4. RESEARCH CHALLENGES IN ACCESS PRICING

For more than a decade, researchers have proposed several static and dynamic access pricing plans for data networks (discussed in Sections 5 and 6). Although most of these remained as models of theoretical interest until 2010, the growing problem of network congestion has created the possibility for realizing new pricing models and schemes in practice. This evolving landscape of access pricing is driven by moves from both wired and wireless ISPs to innovate their data plans as well as by recognition from regulators of the need for such a policy change. Consequently, this renewed interest in pricing has ushered in a new wave of research on SDP [SDP Forum 2012], as evidenced by the growing number of conference tracks and workshops dedicated to this topic [SDP Forum 2012; W-PI 2012; NetEcon 2013]. But research in access pricing is not just about revisiting old ideas, because the Internet ecosystem has evolved rapidly over the past decade. In particular, pricing research today has to account for several factors that have emerged in recent times, such as (a) the inherent time elasticity of demand of several mobile applications, (b) growth in the traffic volume uploaded from the edge devices to the network core, (c) the need for better system scalability and privacy protection in a highly mobile environment, and (d) the ubiquitousness of smart devices with improved graphical user interfaces that enable faster reaction from consumers, thus paving the way to realize demand-response models and other dynamic pricing schemes.

Another key difference is that in addition to appreciating the existing pricing works, researchers need to account for both technological challenges as well as socioeconomic

factors that determine the eventual adoptability and implementability of such pricing plans. Current research on SDP, therefore, focuses on verifying the implementability and efficacy of the proposed models through consumer surveys and field trials of the pricing solutions [Ha et al. 2012a, 2012b; Sen et al. 2013; Dyaberi et al. 2012]. Future research will likely see more such holistic and *multidisciplinary* undertakings between researchers and professionals from engineering, network economics, information systems management, human-computer interaction, marketing and, industry [SDP Forum 2012]. In the next section, we outline several design challenges, open questions, and socioeconomic challenges that researchers need to account for in their work.

#### 4.1. Technological Challenges

*4.1.1. Price Computation.* Determining the optimal price to charge for network access requires developing economic models that account for the ISPs' cost structure (i.e., the trade-off between the cost of overprovisioning and the cost of network congestion) in computing the prices that they are ready to offer and that their consumers are willing to accept. In doing so, the model may require estimating users' delay tolerance for different traffic classes and their deferral patterns in response to offered prices. The key challenges to this are choosing the correct form for utility functions, identifying methods to profile usage behavior, and using these to estimate users' price-delay trade-offs, all while keeping the model computationally tractable. Moreover, for scalability reasons, the ISP may need to perform these estimations without monitoring each individual household's demand pattern so as to avoid additional complexity and communication overhead [Joe-Wong et al. 2011].

Another consideration in price determination is the computation of price elasticity of demand. This is of particular interest in the case of dynamic pricing, which requires some implicit estimation of the users' elasticity across different traffic classes. However, the price a user is willing to pay for individual actions and the elasticity of demand for individual applications is often hidden inside the "bundles" that users purchase. Lastly, coupling dynamic pricing with dynamic QoS is another research area that requires considerable exploration.

*4.1.2. Communication with Users.* The mechanism used to inform users about their usage and available prices also gives rise to several nontrivial research questions: (a) finding ways for ISPs to communicate these information to the user, (b) helping users understand the impact of their usage decisions on their data quota, (c) allowing distributed control of usage across multiple devices under a shared data plan, and (d) understanding the dynamics of quota allocation among users and so forth.

The first question on identifying suitable methods of communicating information to end users depends largely on how the pricing scheme works; for example, in a very dynamic pricing plan, usage decisions will need to be fully automated based on some budget set by the user and auction mechanisms to determine how packets are prioritized [MacKie-Mason and Varian 1995], whereas in less dynamic settings, it may be sufficient to use an application interface for communication [Ha et al. 2012b]. For the second question, developing usage prediction algorithms and integrating these in data-monitoring software applications [Princeton EDGE Lab 2012] can help users to better understand the possible usage implications of their actions. As for the latter two questions, researchers have only begun to investigate the full implications of pricing [Sen et al. 2012a] and quota dynamics of the recently deployed shared data plans. In particular, user trials to understand how pricing practices can affect usage behavior and group dynamics among shared data plan users are some interesting directions for future networking research.




Fig. 3. Schematic of a feedback-control loop between users and ISPs.

**4.1.3. Scalability and Functionality Separation.** One of the design goals of pricing solutions should be system scalability, which requires the identification and separation of client-side and server-side functionalities. For example, dynamic pricing for data needs a feedback-control loop between the ISP server, which computes the prices to offer to users, and the users who respond to these prices [MacKie-Mason et al. 1997]. The ISP-side functionality will typically consist of near real-time traffic monitoring, estimating user delay tolerances for various traffic classes, and computing the prices to offer based on aggregate traffic measurements. On the user side, functionalities like the ability to view future prices and usage history can be provided by installing applications on users' mobile devices. Figure 3 provides a schematic showing the components of this loop.

Ensuring pricing scalability also requires developing models and algorithms that allow ISPs to compute future prices based on aggregate congestion levels, rather than monitoring each customer's usage and response. Other open research questions involve the identification of mechanisms to manage shared data caps in a scalable manner.

**4.1.4. User Personalization.** The personalization of pricing plans has recently received some attention from ISPs, such as in app-based pricing and app bundles, discussed in Section 5.10. The key challenge in dealing with user personalization of data plans is that different users generate different traffic usage patterns by application, location, and time. Consequently, identifying the factors on which such plans should be based and measuring these plans' efficiency need more investigation. Moreover, researchers also need to address the issue of how to track and monetize app-based transaction charges in such personalized plans.

In addition to the wide variation in app usage among different users, app-based pricing faces significant technological challenges on its own. For instance, without Deep Packet Inspection (DPI) software, ISPs may find it difficult to identify the traffic corresponding to different apps. Different users, in fact, may at any given time be using different versions of the same app, which can impact the network in different ways. Moreover, some traffic may not be easily classified as belonging to specific apps—for example, clicking on hyperlinks in certain apps may open the link in an external web browser. Such “out-of-app” traffic may confuse the user as to whether that traffic is counted as part of the original app.

**4.1.5. Consumer Privacy and Security.** The design of access pricing solutions should also account for user privacy, particularly in terms of the feedback that ISPs record from their users and the security of the data. For instance, if users transmit their usage statistics or some measure of willingness to pay from client devices to ISP servers, then this data exchange should take place over a secure connection. In order to view such personal information, users might be required to sign in to an account with their ISP for authentication. Such privacy issues are particularly important for pricing plans that involve per-flow pricing based on willingness to pay; these plans are highly customized to each user and thus involve a substantial exchange of personal information.

A potential way to avoid privacy issues is to minimize the need to monitor individual users' usage patterns, such as by offering uniform prices based on aggregate network loads. Considering only aggregate usage data also helps to ensure scalability to an ISP's full customer base. For any form of dynamic pricing, however, researchers need to address the security aspects of recording, storage, and transfer of information from client devices to ISP servers.

*4.1.6. Mobile Platform and Network Support.* Innovating access pricing requires support from software platforms on user devices and an ISP's network infrastructure, particularly for plans that aim to account for users' response to offered prices. But creating a feedback-control loop as shown in Figure 3 requires developing functionalities for client-side software. However, many widely used platforms (e.g., iOS, Android, and Windows) have different levels of platform openness and can hinder the implementation of such functionalities on the client device. The iOS platform for iPhones and iPads has several restrictions: it strictly specifies what kind of applications can run in the background and further prevents any access other than the standard Application Programming Interfaces (APIs). For example, obtaining an individual application's usage and running a pricing app in the background are prohibited. By contrast, the Android and Windows platforms allow these features (e.g., introducing an API to report individual applications' usage to third-party apps).

Wireless ISPs' current billing systems (including 2G, 3G, and 4G) heavily depend on the Remote Authentication Dial In User Service (RADIUS) protocol, which supports centralized Authentication, Authorization, and Accounting (AAA) for users or devices to use a network service [Rigney et al. 2000]. In particular, RADIUS accounting [Taylor 2000] is well suited to support usage-based pricing, because it can keep track of the usage of individual sessions belonging to each user. Interim-update messages to each session can be sent periodically to update the usage information. However, RADIUS accounting lacks support for dynamic pricing plans, which require ToD usage at various timescales (e.g., hourly, 30 mins, or 10 mins).<sup>4</sup> Consequently, several protocols need to be extended to support these new pricing ideas.

Another interesting direction is to create an open API between user devices and an ISP's billing systems. The open API will foster innovations in pricing for both consumers and providers. For example, the user devices connected to the ISP's network can easily fetch their pricing, billing, and usage information from the network, and the ISP can also easily test and deploy new pricing schemes through the standard interface.

## 4.2. Socioeconomic Challenges

*4.2.1. User Convenience.* Consumers typically prefer the simplicity of flat-rate pricing plans [Odlyzko et al. 2012], and hence, the possibility of inciting negative public reaction is a serious concern for ISPs in experimenting with their data plans. But a combination of smart marketing and efforts to educate consumers has in the past overcome such barriers. For similar reasons, consumers usually prefer a static over a dynamic pricing plan because of the uncertainty associated with dynamic pricing. But this too can be overcome if researchers find the right implementation for dynamic plans and easy-to-use graphical user interfaces to communicate the price- and usage-related information. There are several precedents for dynamic pricing in other markets; for example, electricity markets in the United States have implemented both real-time and day-ahead pricing [U.S. Energy Information Administration 2012], which, with

---

<sup>4</sup>Note that interim update messages are sent periodically when a session joins the system, and hence, the time interval for interim updates should be kept low to support sending ToD usage, which may introduce significant control overhead.



the proliferation of smart meters, are proving to be successful business models. Similarly, in the case of road networks, congestion pricing has been accepted by the public and legislators in cities such as Oslo and Rotterdam [Harshman 2001], where residents agree with the need for tolls to decrease peak time traffic. Moreover, the ability to collect tolls with electronic road pricing has enabled even dynamic, congestion-dependent tolls. The successful adoption of such pricing innovations in these markets indicates that similar pricing ideas may be viable for data networks as well.

To innovate pricing plans for data networks, researchers have proposed mechanisms that have different timescales of price adjustments. In the work of Ha et al. [2012a, 2012b], the prices vary in each hour but are computed and shown to users 1 day in advance through an application that runs on their mobile devices (i.e., day-ahead pricing). In contrast, MacKie-Mason and Varian [1995] proposed a Smart Market scheme that requires real-time response, in which a dynamically varying per-packet price is imposed and users have to place “bids” on each packet to send them into the network. Such an auction-based mechanism requires client-side software for fully automated user responses, improved QoS solutions, and upgrades in the ISP’s billing infrastructure. Researchers therefore need to study not only the impact of the timescale of price adjustments on user behavior and convenience but also the design and implementation aspects of deploying new pricing plans.

*4.2.2. User Budgets.* Consumer trials have shown that mobile data customers exhibit different “willingness to pay,” depending on the application under consideration. For instance, Chen and Jana [2013] conducted a study in which participating users experienced four different speed tiers, each about twice the speed of the previous one. Users were separated into two groups; they either received first higher and then lower, or first lower and then higher, speeds. After these experiences, they were asked how much they were willing to pay for each speed tier. Several users reported remarkably similar willingness to pay for different speed tiers, partly because higher speed tiers do not always result in observable gains across the whole range of applications due to the inherent heterogeneity in apps’ bandwidth needs (e.g., streaming experience improves with higher data speeds, but browsing activities, like mailing, stock apps, are less likely to noticeably improve). In fact, a separate trial [Chen and Jana 2013] asked participants to use specific apps at the different speed tiers. Participants’ willingness to pay for slower speeds dropped by a factor of more than two between low-bandwidth stock apps and streaming YouTube videos. Thus, depending on the applications that a person typically uses, his or her willingness to pay for different levels of service may differ significantly from that of other users. In light of this ambiguity, Chen and Jana propose a framework to allocate different speeds to different users, based on their willingness to pay as negotiated through a second-price auction.

Another determinant of users’ different willingnesses to pay for different speed tiers of service is individual psychology. A recent study by Dyaberi et al. [2012] offered users a probabilistic monetary incentive to stop using data for 10-minute intervals throughout artificially designated “peak hours” of the day. The authors found that psychological measures such as “agreeableness” (compassionate vs. cold) and “openness” (curious vs. cautious) were correlated with users’ acceptance of incentives, whereas “neuroticism” (nervous vs. confident) was associated with noncompliance.<sup>5</sup> Interestingly, such psychological measures were found to be more significant than the type of application in

<sup>5</sup>The authors have three hypotheses related to psychological traits of consumers in their work. Agreeableness hypothesis: someone who is more agreeable will generally be more compliant with requests made of them. Neuroticism hypothesis: someone who is more likely to experience significant negative emotions will be less likely to be compliant. Openness hypothesis: someone curious about new possibilities will be more willing to comply with a new possible mode of usage.

this study, perhaps due to the probabilistic nature of the incentives offered. The authors did not, however, explore the resulting gains in network efficiency from an ISP's perspective, which is an interesting direction for future study.

**4.2.3. Ethical Concerns.** Any pricing innovation will also need to overcome arguments related to ethical issues; for example, consumer advocates have argued that dynamic pricing schemes are unfair to low-income groups, as these users are unable to reduce their usage in peak periods because they have very little usage to begin with. On the other hand, proponents of pricing innovations have pointed out that flat-rate pricing effectively forces the subsidization of heavy users by median and low bandwidth users, and that dynamic pricing enhances users' economic efficiency [Sen et al. 2012b]. This is because dynamic pricing creates the option for all consumers to choose how much they are willing to spend by deciding when and how much bandwidth they want to consume [Ha et al. 2012b]. This view is shared by economists, who describe dynamic pricing for the Internet as a natural extension of existing network control feedback mechanisms that bring "users back into the loop and thereby ensure that performance measures are user-oriented" [MacKie-Mason et al. 1997]. Recent studies in electricity markets have further strengthened this view by confirming that 80% to 90% of low-income households actually stand to gain from dynamic pricing of electricity [Faruqui 2010]. Networking researchers should be aware of these ethical arguments regarding access pricing in order to overcome policy hurdles.

### 4.3. Regulatory Challenges

Pricing in data networks has remained a politically charged issue, particularly for pricing mechanisms that can potentially create incentives for price discrimination, nonneutrality, and other anticompetitive behavior through app-based pricing or bundling of access and content. Academics in law schools have already cautioned that the ongoing debate on network neutrality in the United States often overlooks service providers' need for flexibility in exploring different pricing regimes [Yoo 2009]:

*Restricting network providers' ability to experiment with different protocols may also reduce innovation by foreclosing applications and content that depend on a different network architecture and by dampening the price signals needed to stimulate investment in new applications and content.*

But faced with the growing problem of network congestion, there has been a monumental shift in the regulatory perspective in the United States and other parts of the world. This sentiment was highlighted in FCC chairman J. Genachowski's December 1, 2010, statement [Schatz and Ante 2010], which recognizes "the importance of business innovation to promote network investment and efficient use of networks, including measures to match price to cost."

## 5. STATIC PRICING

### 5.1. Fixed Flat-Rate Pricing

**5.1.1. Monthly Rate.** In the past, ISPs have charged users a flat monthly fee for broadband access, irrespective of the actual time spent on the network or data usage. Such a simple pricing model was praised by Anania and Solomon [1997] and Odlyzko [2001] based on historical and political precedence. A more theoretical argument was made by Shakkotai et al. [2008], showing that simple flat-rate pricing is quite efficient in extracting revenue for elastic traffic if all users run an identical application, but with different valuations for the application.

There are several variations of the monthly flat-rate pricing model in the real world. For example, "unlimited" data plans put no cap on the bandwidth used every month. If a maximum usage limit is predetermined according to a flat price, the plan is called

“flat up to a cap.” Exceeding the limit usually incurs penalty costs that are proportional to the usage above the cap (i.e., “metered” pricing). “Tiered” data plans with different flat-rate pricing for different usage caps are often used to provide a range of choices to consumers. Usually, users on flat-up-to-a-cap plans are shifted to usage-based pricing or a higher-priced tier upon exceeding their data limit. But another emerging trend is “flat to a cap, then throttle,” as offered by Orange Spain and Comcast [Sephton 2011b; Comcast 2012].

Without such penalties, flat-rate plans will simply become unviable for ISPs. As if to drive home that point, TelstraClear, New Zealand’s second largest telecommunications provider, experimented with a free weekend plan during which they switched off their usage metering and removed data caps from Friday evening to midnight on Sunday from December 2–5, 2011. When all their data-hungry customers simultaneously descended on the net with whetted appetite for the big weekend buffet, many were dismayed to find their speeds down to one-fifth of the usual [Keall 2011].

*5.1.2. Hourly Rate.* Some providers have begun offering flat-rate plans in which mobile Internet services are billed by the hour—that is, the cap is specified in terms of time instead of usage. For example, USB modem customers of the Egyptian mobile operator Mobinil can choose from packages of 30 hours for EGP 80/month or 60 hours for EGP 125/month [Telecompaper 2010]. Time-based pricing has also been proposed for road networks, in which users are charged based on how long it takes to travel from one location to another [May and Milne 2000]. However, a London study showed that time-based pricing encourages reckless driving by rewarding users who drive faster [Richards et al. 1996].

Even in broadband networks, hourly and other flat-rate pricing plans have significant disadvantages. Although flat-rate billing is cheap to implement and operate, encourages user demand, and creates simple and predictable monthly fees for customers, it suffers from several disadvantages. First, flat-rate plans lead to inefficient resource allocation and market segmentation, with low-usage customers typically subsidizing the heavy users (i.e., the bandwidth hogs) [Hendershott 2006]. Second, while ISP revenues depend on the median user, peak load costs are driven by the heavy users, which creates a price-cost mismatch. In the Berkeley INDEX project, Varaiya [1999] concurred with many of these disadvantages. Consequently, as the demand for bandwidth grows, most ISPs are replacing their flat-rate plans with usage-based or metered data plans.

## 5.2. Usage-Based Pricing

New Zealand conducted an early experiment with usage-based pricing in 1989, when six participating universities agreed to pay by the volume of traffic sent in either direction through the NZGate gateway [Brownlee 1997]. Such a pricing plan was necessitated by New Zealand’s geographically isolated location, expensive trans-Pacific links, and the lack of subsidies from the New Zealand government [McKnight and Bailey 1997]. This arrangement provided an early demonstration that metering traffic and charging users by daily volume is feasible.

“Metered” pricing implies that a user is charged in proportion to his or her actual volume of data usage [Walrand 2008; Li et al. 2009]. In practice, operators often use “cap then metered” plans (also known as usage-based pricing), for which a user pays a flat price up to a predetermined volume of traffic, beyond which the user is charged in proportion to the volume of data consumed [Hande et al. 2010]. These variations on usage-based pricing are the prevalent pricing plans in many countries, including the United States, United Kingdom, Korea, Japan, and Hong Kong [Kang 2010; Parker 2010; Morgan 2011].

Similar pricing schemes have also been implemented for other networks. For instance, analogous *distance-traveled* pricing, in which drivers are charged based on the distance they travel, has been implemented for vehicle traffic in the United States [Holguin-Veras et al. 2006], as well as for trucks in Switzerland, Austria, and Germany [McKinnon 2006].

### 5.3. Paris Metro Pricing

Paris Metro Pricing (PMP) was proposed by Odlyzko 1999 as a simple and elegant solution for creating differentiated service classes. PMP partitions the network resources into several logical traffic classes, each of which is identical in its treatment of data packets but charges users differently. Thus, users willing to pay more will select the more expensive, and hence less congested, logical traffic class.

PMP is designed to enable maximum simplicity for the end user in expressing his or her preference through self-selection of the desired service level. But Odlyzko [1999] also identified some potential problems in implementing PMP, such as finding ways to set the prices and capacities of the logically separate channels, and to maintain predictable performance on the different channels to avoid network instability. PMP may also require better designs for the user interface to let users dynamically alter their preferences and assign application sessions to different traffic classes. Recently, Wahlmueller et al. [2012] proposed a more user-centric variant on PMP in which QoS classes are replaced by QoE classes, with parameters such as bandwidth rate and packet loss adjusted so that the user's quality of experience remains at a specified level for different types of applications (e.g., streaming requires higher bandwidth than web browsing to achieve the same QoE). Users can also change their desired QoE, and the corresponding price they are charged [Reichl 2010].

Despite the deployment challenges for broadband networks, PMP takes its name from an actual pricing scheme on the Paris metro that was used until the 1980s to differentiate first- and second-class coaches. A similar concept is used today in the United States for roads with High-Occupancy Vehicle (HOV) lanes, which are identical to all other lanes of the road, except that HOV lanes are restricted to vehicles with multiple passengers [Poole and Orski 2003]. Drivers can self-select to use less-congested HOV lanes by paying the higher "price" of carpooling with other passengers.

### 5.4. Token Pricing

Lee et al. 2011 introduced a token pricing scheme in which users pay a fixed flat-rate monthly fee for Internet access. Each user receives a certain number of "tokens" from the service provider, which offers two service classes of different qualities. The basic-quality class requires no tokens for access but may become congested. The higher-quality class requires users to redeem some tokens and hence has lesser congestion and a better service because of a PMP-like effect. Thus, users self-prioritize their sessions and implicitly pay for the higher QoS with their tokens for particularly urgent sessions during peak network congestion. Each session, regardless of size, costs the same number of tokens, and users continually receive more tokens (e.g., one every 10 minutes). The benefits of such a system in reducing peak network congestion in a real network have yet to be explored.

### 5.5. Priority Pricing

Cocchi et al. [1991, 1993] studied a pricing scheme in a multiple service network with priority classes but without any resource reservation. Users can request different QoS by setting bits in their packets. A higher-priority class charges a higher per-byte fee but is assumed to receive better service from the network. Thus, users who pay a greater per-byte fee for higher priority are in effect paying for the negative externality

imposed on traffic from other, lower-priority users. The authors showed that such quality-sensitive pricing is more efficient (in a Pareto sense) than flat pricing. However, this result depends on the reservation-less assumption.

A game-theoretic framework was used to analyze static priority pricing by Marbach [2004]. In his single-link model, users assign a priority class to their packets and are charged accordingly. These charges are based on the packets submitted to the network rather than their actual delivery. Marbach showed that there always exists a Wardrop equilibrium bandwidth allocation, but it is not necessarily unique. This equilibrium allocation and the associated link revenue do not depend on the prices of the different priority classes, but the prices do have a simple relation with the packet loss probability due to congestion in that class.

One disadvantage of priority pricing is users' inability to express their desired levels of delay and bandwidth share; priority pricing's consistent preferential treatment of higher-priority classes might drive lower-priority classes to little or no usage. Priority pricing is largely absent today, but some ISPs are considering the idea of creating a "priority data plan" in which users of the premium service are prioritized during periods of network congestion. Recently, SingTel of Singapore introduced such an option, called "priority pass," for its top-tier dongle customers [Kwang 2011].

## 5.6. Reservation-Based Pricing

Parris et al. [1992] were one of the first to study pricing in a reservation-oriented network. They considered the issues of network utilization, ISP revenue, and blocking probability under per-packet, setup, and peak-load pricing schemes. In their work, users are characterized by their connection durations, budgets, and chosen classes of service (with a higher per-byte fee for a higher-priority service class). The network decides to either accept or block the connection, depending on the sufficiency of the user's budget and the availability of network resources. Using simulations, the authors show that for a given pricing scheme, price increases will at first increase and then eventually decrease the net revenue, but will always decrease the blocking probability and network utilization. Setup pricing decreases the blocking probability and increases revenue for the ISP, and more generally performs better than per-packet pricing (in that the blocking probability from admission control is lower under setup pricing than per-packet pricing for the same level of revenue generated).

Despite these advantages, Parris et al.'s [1992] form of reservation pricing suffers from some shortcomings. First, a flat-rate setup cost is unfair toward those users with shorter conversations. Second, poorer users may not be able to afford a connection under a high setup cost, thus leading to a greater digital divide. Third, average network utilization is typically lower in the presence of setup costs, thus demonstrating a trade-off between network efficiency and revenue maximization.

Parris and Ferrari [1992] presented another reservation pricing scheme for real-time channel establishment. Under this plan, users are charged based on the type of service requested, as measured by factors such as the bandwidth, buffer space, CPU time resources reserved, and delay imposed on other users. The total charge a user pays is a product of the type of service measure, channel duration, and a ToD factor. However, the work does not provide any clear guidelines on how economic considerations are to be mapped to a single ToD factor or the impact of the overhead associated with estimating network parameters in real time.

In a later work, Delgrossi and Ferrari [1999] considered a pricing scheme based on the portion of resource capacity used by a reserved data channel in a multiple-service network. They introduced a charging formula with different reservation and transport cost components for real-time and non-real-time traffic, along with a discussion on



computing resource capacity requirements for the channel as a function of the buffer size, processing power, and schedulability.

### 5.7. Time-of-Day Pricing

ToD pricing schemes charge peak and off-peak hours differently, so as to disperse user demand more uniformly over time. Parris et al. [1992] and Parris and Ferrari [1992] considered a form of ToD combined with reservation-based pricing, which divides a day into peak and off-peak periods and incorporates the time elasticity of user demand. They show that peak-load pricing reduces the peak utilization and the blocking probability of all traffic classes and increases revenue by inducing a more even distribution of demand over peak and off-peak periods.

The most basic form of ToD in practice is a two-period plan that charges different rates during the daytime and nighttime. For example, BSNL in India offers unlimited nighttime (2–8 am) downloads on monthly data plans of Rs 500 (\$10) and higher. Other variations of ToD pricing are offered elsewhere; for instance, the European operator Orange has a Dolphin Plan for £15 (\$23.50 USD) per month that allows unlimited web access during a “happy hour” corresponding to users’ morning commute (8–9 am), lunch break (12–1 pm), late afternoon break (4–5 pm), or late night (10–11 pm).

ToD pricing has been implemented in road networks since the 1990s in California, Norway, and Sweden [Gomez-Ibanez and Small 1994]. The electricity market has also practiced static ToD pricing for many years, a move accelerated by power shortages over the past decade. Many academic works on ToD pricing in these markets have also focused on constructing user demand functions for ToD pricing by using data on electricity consumption, focusing on the case of two prices (peak and off-peak) [Faruqui and Wood 2008; Hausman et al. 1979]. Several works have reported on consumer trials of ToD pricing [Charles River Associates 2005; Faruqui et al. 2009; Herter 2007; Matsukawa 2001; IBM 2007; Wells and Haas 2004; Wolak 2006]. Many of these trials employed a variant of ToD, called *critical peak pricing*, in which the price on certain days of especially high electricity consumption (e.g., very hot days during the summer) was increased beyond the normal peak price. Such pricing was generally judged to be successful in lowering peak electricity consumption, especially on days of high demand.

### 5.8. Expected Capacity Pricing

In 1997, Clark wrote:

*In the future it will be desirable to provide additional explicit mechanisms to allow users to specify different service needs, with the presumption that they will be differentially priced.*

He proposed expected capacity pricing as a mechanism to allow users to explicitly specify their service expectation (e.g., file transfer time) while accounting for differences in applications’ data volume and delay tolerance. The idea is that by entering into profile contracts for expected capacity with the operator, different users should receive different shares of network resources *only* at times of congestion [Songhurst 1999].

One specific proposal to realize this service involved traffic flagging (i.e., each packet is marked as being *in* or *out* of the user’s purchased profile, irrespective of network congestion level) by a traffic meter at access points where the user’s traffic enters the network. This is followed by congestion management at the switches and routers where packets marked as *out* are preferentially dropped during congested periods, but are treated in an equal best-effort manner at all other times. The expected capacity is thus not a capacity guarantee from the network to the user, but rather a notion of the capacity that a user expects to be available and a set of mechanisms that allow him or her to obtain a different share of the resource at congested times. This pricing can be simply enforced at the router and switches of the network, and allows service

providers to have more stable estimates of the future necessary capacity based on the total expected capacity sold, rather than the sum of peak rates of all users' access links. A dynamic pricing version of the scheme has also been explored [Clark 1997]. However, in order to implement this pricing scheme, the assignment of price values to expected capacity profiles requires further study.

### 5.9. Cumulus Pricing

Cumulus pricing schemes consist of three stages: specification, monitoring, and negotiation. A service provider initially offers a flat-rate contract to the user for a specified period based on the user's estimate of resource requirements. During this time, the provider monitors the user's actual usage and offers periodic feedback to the user (by reporting on "cumulus points" accumulated from their usage) to indicate whether the user has exceeded the specified resource requirements. Once the cumulative score of a user exceeds a predefined threshold, the contract is renegotiated. Hayel and Tuffin [2005] studied such a scheme and used simulated annealing to optimize the total network revenue in terms of the renegotiation threshold.

Cumulus pricing is a simple pricing scheme that can be easily implemented at the network edge. Some ISPs have been experimenting with similar ideas. For example, Vodafone in the United Kingdom announced a new "data test drive" plan that allows customers joining any of the monthly pay contracts to have unlimited data access (including tethering, but excluding roaming) for the first 3 months. The data usage report is then fed back to the user to negotiate whether his or her chosen plan is appropriate. The user can choose to either continue with the existing plan, possibly incurring overage charges, or to switch to an alternative plan suggested by Vodafone [Sephton 2011c].

### 5.10. Application- and Content-Based Pricing

Several mobile service providers have been experimenting with various forms of application- or content-based pricing. These are somewhat analogous to certain types of tolls in road networks, in which drivers are charged different rates for different types of vehicles (e.g., trucks are charged at a higher rate) [U.S. Office of Highway Policy Information 2011]. Indeed, charging by vehicle type was mentioned as a criterion for efficient road pricing in a report by the U.K. Ministry of Transportation [May 1992].

Although the currently available app-based data plans are mostly designed to attract and "lock in" customers by bundling content and data plans, they highlight an emerging trend of operators experimenting with pricing structures that charge (or subsidize) differently based on application type [Ericsson 2011; Higginbotham 2010]. These types of bundled services, which generally offer bundled access to music or movie streaming, are offered by Three in the United Kingdom, Telus in Canada, and TDC in Denmark [Sephton 2011a; Vardy 2011]. Orange UK allows users to choose which media services to bundle [Owoseje 2011].

Another form of app-based pricing that is becoming increasingly popular is "sponsored content," in which operators subsidize certain types of data. For instance, Mobistar in Belgium offers "zero-rated" or free access to Facebook and Twitter [Bubley 2011]. Japan's Softbank Mobile is also moving in this direction, having contracted with Ericsson to develop the infrastructure for app-based and other types of innovative data plans [Ericsson 2010]. Whereas technological advances such as DPI have allowed ISPs to differentiate between the traffic of some apps and offer these sponsored content plans, app-based pricing still faces several challenges. As discussed in Section 3, these challenges will likely limit the deployment of app-based pricing at the present time.

Figure 4 compares the various static pricing schemes discussed in this section in terms of their key features, deployment status, and both technological and

|   | Usage Based  | Paris Metro  | Token   | Priority   | Reservation  | Time-of-Day  | Expected Capacity   | Cumulus   | App Based  |
|---|--|--|---|--|--|--|---|---|--|
| <b>Model features</b>                       | <ul style="list-style-type: none"> <li>• Users pay by volume of data.</li> <li>• Can include data caps, over which users may be throttled or pay overages.</li> </ul>                        | <ul style="list-style-type: none"> <li>• Two QoS classes with different rates.</li> <li>• Users self-select so as to ensure promised QoS.</li> </ul> | <ul style="list-style-type: none"> <li>• Flat access fee for basic QoS.</li> <li>• Users receive tokens to be redeemed for higher QoS.</li> </ul> | <ul style="list-style-type: none"> <li>• Users request higher QoS for certain packets.</li> <li>• Higher rates are charged for better QoS.</li> </ul>    | <ul style="list-style-type: none"> <li>• Network admission based on user budget and available network resources.</li> </ul>    | <ul style="list-style-type: none"> <li>• Different usage-based rates for “peak” and off-peak hours.</li> </ul>                             | <ul style="list-style-type: none"> <li>• Users contract with provider for service expectations in advance.</li> </ul>                             | <ul style="list-style-type: none"> <li>• Service provider offers initial contract.</li> <li>• Renegotiation at usage-based “cumulus points.”</li> </ul> | <ul style="list-style-type: none"> <li>• Different apps charged by different usage-based rates.</li> <li>• “Zero-rated” apps with free traffic.</li> </ul> |
| <b>Deployment status</b>                    | <ul style="list-style-type: none"> <li>• Most UK, Korean ISPs</li> <li>• NTT DoCoMo (Japan)</li> <li>• CSL (Hong Kong)</li> <li>• AT&amp;T, Verizon (US)</li> <li>• Road networks</li> </ul> | <ul style="list-style-type: none"> <li>• Simulation</li> <li>• HoV lanes in road networks</li> </ul>   | <ul style="list-style-type: none"> <li>• Simulations</li> </ul>   | <ul style="list-style-type: none"> <li>• SingTel (Singapore)</li> </ul>  | <ul style="list-style-type: none"> <li>• Simulations</li> </ul>  | <ul style="list-style-type: none"> <li>• BSNL (India)</li> <li>• Orange (Europe)</li> <li>• Electricity market</li> </ul>                  | <ul style="list-style-type: none"> <li>• Proposal</li> </ul>  | <ul style="list-style-type: none"> <li>• Vodafone (UK)</li> </ul>   | <ul style="list-style-type: none"> <li>• Mobistar (Belgium)</li> <li>• Telus (Canada)</li> <li>• TDC (Denmark)</li> <li>• Three (UK)</li> </ul>            |
| <b>Technological challenges</b>             | <ul style="list-style-type: none"> <li>• Monitoring usage amounts.</li> </ul>  | <ul style="list-style-type: none"> <li>• Calculating prices of different QoS classes.</li> <li>• Setting capacities of different classes.</li> </ul> | <ul style="list-style-type: none"> <li>• Tracking users’ token budgets.</li> <li>• Calculating token accumulation rates.</li> </ul>               | <ul style="list-style-type: none"> <li>• Calculating prices for different priorities.</li> <li>• Ensuring a minimum QoS for lower priorities.</li> </ul> | <ul style="list-style-type: none"> <li>• Keeping track of user budgets.</li> <li>• Monitoring of network resources.</li> </ul> | <ul style="list-style-type: none"> <li>• Choosing hours to designate as “peak.”</li> <li>• Monitoring usage at different times.</li> </ul> | <ul style="list-style-type: none"> <li>• Negotiating contracts with each user.</li> <li>• Assigning price values to expected capacity.</li> </ul> | <ul style="list-style-type: none"> <li>• Monitoring of cumulus points.</li> <li>• Calculating initial contract offers.</li> </ul>                       | <ul style="list-style-type: none"> <li>• Monitoring traffic of multiple distinct apps.</li> <li>• Classifying inter-app traffic.</li> </ul>                |
| <b>Necessary features for user adoption</b> | <ul style="list-style-type: none"> <li>• Interface for users to track their consumption.</li> </ul>  | <ul style="list-style-type: none"> <li>• Interface to show QoS received.</li> <li>• Guidance in choosing between QoS classes.</li> </ul>             | <ul style="list-style-type: none"> <li>• Interface to keep track of tokens.</li> <li>• Automated guidance on when to use tokens.</li> </ul>       | <ul style="list-style-type: none"> <li>• Automated algorithm to choose priorities.</li> <li>• Interface to request higher QoS.</li> </ul>                | <ul style="list-style-type: none"> <li>• Interface to show admitted sessions and user budget.</li> </ul>                       | <ul style="list-style-type: none"> <li>• Interface to track usage at different times of the day.</li> </ul>                                | <ul style="list-style-type: none"> <li>• Automatic contract negotiation.</li> <li>• Guidance on service specification.</li> </ul>                 | <ul style="list-style-type: none"> <li>• Interface to track cumulus points.</li> <li>• Automated contract negotiation.</li> </ul>                       | <ul style="list-style-type: none"> <li>• Interface to track usage by app.</li> <li>• Clear distinctions between apps.</li> </ul>                           |

Fig. 4. Comparison of static pricing schemes (Section 5).

socioeconomic challenges. Although different forms of usage-based pricing are currently the norm in most parts of the world, some operators have begun to introduce more innovative data plans, such as priority- or app-based pricing.

## 6. DYNAMIC PRICING

In this section, we discuss pricing models in which the prices offered vary over time (e.g., due to fluctuating network congestion levels). Although such models often allow “better” network management by providing the ISP with more pricing flexibility, they also suffer from significant barriers to user adoption, as discussed in Section 4. A summary and comparison of the pricing plans discussed may be found in Figure 5.

### 6.1. Raffle-Based Pricing

Like ToD pricing, raffle-based pricing divides the day into peak and off-peak periods. Users are then offered probabilistic incentives to shift their demand to off-peak periods, in the form of a raffle or lottery for a given monetary reward. The probability of winning the lottery is proportional to the user’s contribution to the reduction in peak demand. This all-or-nothing lottery may instead be replaced by one in which the ISP divides the total reward by the total amount of traffic shifted, paying this amount to each user with a probability equal to the percentage of usage shifted to the off-peak period. Loiseau et al. [2011] derived expressions for the Nash equilibrium of this user-ISP interaction and show that in some cases the social welfare is more robust to price variations than a comparable ToD pricing plan with two periods. Yet the uncertainty inherent in raffle-based pricing may reduce its effectiveness; users may not shift their usage to off-peak periods without guaranteed rewards. Indeed, the reward amount depends on other users’ behavior. A field trial of such a policy would therefore be necessary to understand its effectiveness.

|                                      | Raffle Based   | Priority Pricing   | Effective Bandwidth   | Auctions (Smart Market)  | Congestion Pricing  | Day Ahead   |
|--------------------------------------|--|--|---|--|---|---|
| Model features                       | <ul style="list-style-type: none"> <li>Lottery to reward users for shifting traffic to off-peak times.</li> <li>Each user's probability of winning equals her contribution to the total amount shifted.</li> </ul> | <ul style="list-style-type: none"> <li>Users pay higher prices for more immediate service.</li> <li>Prices change with estimates of users' preferred waiting times.</li> </ul> | <ul style="list-style-type: none"> <li>Charge by user-predicted peak and mean traffic rates and session duration.</li> <li>Flat fee to renegotiate the session's cost.</li> </ul> | <ul style="list-style-type: none"> <li>Users submit bids to send each packet through the network.</li> <li>ISP admits packets with high bids and charges users the lowest admitted bid.</li> </ul> | <ul style="list-style-type: none"> <li>Price changes based on real-time network congestion.</li> <li>Users adjust their usage amounts based on the current prices.</li> </ul> | <ul style="list-style-type: none"> <li>Guarantee hourly prices 1 day in advance.</li> <li>Dynamically re-compute prices each day based on observed usage changes due to previous prices.</li> </ul> |
| Deployment status                    | Simulations  | Simulations  | Simulations   | Simulations  | Simulations   | Pilot trial   |
| Technological challenges             | <ul style="list-style-type: none"> <li>Calculating amount of traffic shifted per user.</li> <li>Raffle to choose the winning user.</li> </ul>  | <ul style="list-style-type: none"> <li>Modifying topology management, databases, and accounting.</li> </ul>  | <ul style="list-style-type: none"> <li>Modifying accounting system.</li> <li>Calculating effective bandwidth.</li> </ul>  | <ul style="list-style-type: none"> <li>Coordination among different network nodes.</li> <li>Modifying accounting system.</li> </ul>  | <ul style="list-style-type: none"> <li>Real-time measurement of network load.</li> <li>Calculating prices.</li> </ul>   | <ul style="list-style-type: none"> <li>Communicating day-ahead prices to users.</li> <li>Modifying billing system.</li> </ul>   |
| Necessary features for user adoption | <ul style="list-style-type: none"> <li>User uncertainty due to lottery.</li> <li>Interface to show probability of winning.</li> </ul>  | <ul style="list-style-type: none"> <li>Automated algorithms to select waiting times.</li> <li>Interface to show delayed sessions.</li> </ul>                                   | <ul style="list-style-type: none"> <li>Automated algorithms to predict session bandwidth rates.</li> <li>Interface to show connection charges.</li> </ul>                         | <ul style="list-style-type: none"> <li>Automated bidding algorithms.</li> <li>Interface to show accumulated charges and packet delay.</li> </ul>   | <ul style="list-style-type: none"> <li>Automated mechanism to respond to fluctuating prices.</li> <li>Real-time transmission of prices to users.</li> </ul>                   | <ul style="list-style-type: none"> <li>Interface to show day-ahead prices.</li> <li>Automated algorithm to schedule usage.</li> </ul>   |

Fig. 5. Comparison of dynamic pricing schemes (Section 6).

A recent trial by Dyaberi et al. [2012] employed a similar probabilistic incentive mechanism, offering (with some probability) rewards to users for terminating their cellular usage during a 10-minute interval. Although the authors observed an 83% acceptance rate, their trial participants were largely university students, and the devices used were loaned for the purposes of the trial. Thus, their results may not extend to the general population.

## 6.2. Dynamic Priority Pricing

Gupta et al. [1997] presented a dynamic priority pricing mechanism in which the prices serve as a congestion toll for network access. The authors modeled user service requests as a stochastic process and network nodes as priority queues. A user's incoming request has an instantaneous value for the service and a linear rate of decay for this value to capture the delay cost. Users' requests can be fulfilled using different alternatives, each with its corresponding price and waiting time. The user trades off between the total cost of service and the cost of delay to choose an optimal alternative in a particular priority class. The principal benefit of such a scheme is that it achieves resource allocation in real time in a completely decentralized manner; moreover, users' myopic decisions lead to socially optimal choices.

Dynamic priority pricing is based on general equilibrium theory, which generalizes the notion of an "equilibrium" to several interacting markets—in this case, the market situation at any particular point in time (i.e., supply equals demand at any given time). However, since computing general Arrow-Debreu equilibria in the volatile environment of the Internet is expensive, the authors introduced the notion of a stochastic equilibrium and derive optimal prices that support the unique stochastic equilibrium. They also developed a decentralized real-time mechanism to compute near-optimal prices for the stochastic equilibrium: the cost of service (i.e., the price) is updated after each time period by taking a weighted average of the price in the previous period and the new optimal price, calculated from updated estimates of the waiting times for different priority classes. The convergence properties of this dynamic pricing mechanism were demonstrated with simulations that allowed the system to adapt to changing loads on

the network. However, implementing dynamic priority pricing will require modifications to network topology management applications, databases, accounting systems, and the end-user interface.

### 6.3. Proportional Fairness Pricing

Kelly et al. [1998] proposed proportional fairness pricing as a means to allocate resources (which determine user rates) in proportion to the user's willingness to pay. The global optimization of maximizing net utility across all users, given resource capacity constraints, can be decomposed into a user and a network optimization problem. Kelly showed that there exist a price vector and a rate vector that optimize both the user and the network's optimization problems. Alternatively, if each user chooses a price per unit time according to his or her willingness to pay, and if the network allocates rates per unit price that are proportionally fair, then a system optimum is achieved when users' choices of prices and the network's rate allocation are in equilibrium. Courcoubetis et al. [1998] extended this idea by replacing end users with intelligent agents that can decide the willingness to pay on behalf of the user while maximizing the user's utility. However, this element introduces overhead in installing such agents on the users' devices or machines and adding network servers to compute the optimal rate allocation vectors at a short timescale.

### 6.4. Effective Bandwidth Pricing

Kelly's [1994] effective bandwidth pricing is a variant of usage-based pricing in which users are charged based on self-reported peak and mean traffic rates, as well as the observed mean rate and duration of each connection. Before a user's connection is accepted, the user is required to provide mean and peak rates for the connection. Given a formula describing effective bandwidth as a function of the peak and mean rates, the user is charged a tariff given by the tangent line to this effective bandwidth formula (as a function of the mean rate) at the self-reported mean and peak rates. Evaluating this tariff at the observed mean rate, the result is multiplied by the connection duration to give the total charge to the user. Figure 6 shows a schematic illustration of such a tariff.

Kelly [1994] showed that under this pricing scheme, users minimize their expected cost by accurately reporting the connection's mean and peak rates. Thus, the final charge to the user consists of a term proportional to the connection duration and another term proportional to the connection volume. Users may renegotiate the tariff for a flat fee, such as for highly variable traffic.

This pricing scheme can also be extended to connection acceptance control, in which case, a connection is accepted if the network's effective load, as calculated from the tariffs charged to existing connections, falls below a certain threshold value. Although effective bandwidth pricing is compatible with user incentives and fairly simple, it does require an explicit effective bandwidth formula, and it requires users to know, or at least estimate, the peak and mean rates of each connection. Moreover, further validation is needed to understand whether the benefits of this pricing scheme justify the accounting overhead associated with charging each connection based on its duration and volume.

### 6.5. Responsive Pricing

MacKie-Mason et al. [1997] described the concept for responsive pricing in the following words:

*By associating a cost measure with network loading, all users can be signaled with the prices necessary to recover the cost of the current network load. Price-sensitive users—those willing and able to respond*






Fig. 6. Illustration of the tariff charged in effective bandwidth pricing [Kelly 1994].

*to dynamic prices—increase economic efficiency by choosing whether or not to input traffic according to their individual willingness to pay the current price.*

In other words, a user’s price sensitivity and time sensitivity for different applications can be exploited by networks to dynamically set prices to alleviate congestion. This process broadly encompasses the philosophy behind different forms of dynamic time-dependent pricing. In the case of the Internet, MacKie-Mason et al. [1997] argued that such a responsive pricing system is required for network efficiency.

The network can set response prices either in a closed-loop feedback [Murphy et al. 1994; Murphy and Murphy 1994] or a Smart Market approach [MacKie-Mason and Varian 1995]. In a closed-loop setting, the network state, measured in terms of the buffer occupancy at the gateway, is converted to a price per packet for users’ adaptive applications, which then decide how much data to transmit. This closed-loop feedback is thus similar to the priority pricing of Gupta et al. [1997], discussed in Section 6.2: the basic idea is that dynamic prices are set so as to incentivize users to behave in a way that enhances network efficiency. However, due to the adaptive nature of the prices, there is a slight delay in the feedback loop: prices are set based on network conditions in the previous time period.<sup>6</sup>

This delay in the feedback loop is tightened in the Smart Market approach proposed by MacKie-Mason and Varian [1995]. In this approach, each user places a “bid” on each packet that reflects their willingness to pay to send the packet onto the network at a given time. The gateway admits packets in the descending order of their bids as long as the network performance remains above a desired threshold. Users are charged according to the minimum bid on a packet admitted into the network at the time, and

<sup>6</sup>Similar ideas have also been adopted in the electricity industry; for instance, Vytelingum et al. [2010] developed an auction scheme that takes into account varying market capacity, with dynamic offers from electricity distributors and real-time responses from users (i.e., households) buying electricity.

thus they pay only for the congestion cost at the market clearing price. Although such auction schemes encourage network and economic efficiency, they require substantial changes in users' application interfaces and providers' billing systems, with additional concerns in the case of billing contention and so forth. Moreover, as discussed by Sarkar [1997], such a pricing scheme would require extensive coordination among different nodes of the network, and is open to abuse by those controlling bottleneck facilities. Sarkar argued that these concerns may be overcome with appropriate governmental regulation.

### 6.6. Dynamic Congestion Pricing

Dynamic congestion pricing is a particular realization of the idea of responsive pricing, in which the network announces prices based on current congestion levels and the user response to these prices is fed back into the control loop to compute new prices. Ganesh et al. [2001] used congestion prices as a mechanism to provide both feedback and incentives to end systems for rate adaptation in a decentralized manner; they then studied the resulting dynamics. Paschalidis and Tsitsikilis [1998] addressed the issue of revenue and welfare maximization for dynamic congestion pricing of customer calls by using a dynamic programming formulation. In their model, users initiate calls that differ in their service class, resource requirements, and call duration. Based on the current congestion level, the service provider charges a fee per call, which in turn influences users' demand. Their findings additionally corroborate the usefulness of ToD pricing in reducing network congestion problems.

Congestion-based pricing innovations have been adopted in many markets outside the United States. In recent years, network operators in highly competitive and lucrative markets, such as those in India and Africa, have adopted innovative congestion- and location-dependent dynamic pricing for voice calls [The Economist 2009; Sen 2011], although not yet for mobile data plans. An analogous congestion-based scheme for transportation networks, in which users are charged based on both their distance traveled and time to travel that distance, was considered in Cambridge, U.K., but was never fully implemented [Gomez-Ibanez and Small 1994].

Congestion pricing in wireless cellular networks has also received some attention in the literature, such as in a survey by Al-Manthari et al. [2011], which divides pricing policies into two categories: admission-level and power-level pricing. The authors mentioned several challenges specific to wireless networks (e.g., interference, user mobility, and bandwidth limits) and surveyed proposed pricing policies that incorporate priority queues and QoS support.

Several works on pricing in the electricity market have adopted a congestion-based dynamic pricing model, as reviewed by Borenstein et al. [2002] and proposed in a more recent paper by Samadi et al. [2010]. Many of these focus on end-user management of congestion-dependent prices, proposing algorithms to predict future prices and schedule energy usage accordingly [Du and Lu 2011; Mohsenian-Rad and Leon-Garcia 2010; Mohsenian-Rad et al. 2010]. Other works have attempted to combine these two perspectives by examining a feedback loop between the user and provider [Roosbehani et al. 2010; Li et al. 2011].

### 6.7. Day-Ahead Pricing

Day-ahead pricing is a variant of dynamic congestion pricing in which prices vary by the time of the day (e.g., each hour) and users are told the future prices 1 day in advance. Joe-Wong et al. [2011] proposed an algorithm for calculating these day-ahead prices that models a user's willingness to shift different types of usage to lower-price times. The model accounts for the fact that some types of traffic (e.g., email) are more delay sensitive, whereas for other traffic (e.g., software updates), users may be willing

to wait several hours for a lower price. As prices are offered throughout the day, the ISP monitors the change in traffic volume relative to a baseline traffic trace with constant prices. These changes are then used to optimize the prices over the next day, in a feedback loop structure similar to that proposed by MacKie-Mason et al. [1997].

A recent pilot trial in the United States has demonstrated the feasibility, benefits, and architecture of day-ahead pricing for mobile data. In this trial, Ha et al. [2012b] implemented a form of day-ahead pricing for a small number of iOS users and reported that such pricing policies can in fact reduce the peak load on the network. Although this form of dynamic pricing less accurately reflects the real-time network congestion, it has the advantage of increasing user certainty by guaranteeing some prices in advance and allowing users to plan ahead. In fact, the trial of Ha et al. [2012b] included an automated scheduling algorithm on user devices that can schedule usage based on future prices so as to save users' money. A similar day-ahead pricing scheme has also been proposed for electricity smart grids [Joe-Wong et al. 2012a].

### 6.8. Game-Theoretic Pricing

Several authors have used game-theoretic models for pricing data, some of which are briefly discussed here. Hayer [1993] proposed transport auctions as a way to distribute excess capacity across users with delay-tolerant traffic. A decentralized auction-based approach to the pricing of edge-allocated bandwidth, called "market pricing," was explored by Semret et al. [2000] in a differentiated services Internet.

Yaïche et al. [2000] introduced a cooperative game-theory framework that used Nash bargaining (i.e., a competition model whose solution satisfies certain fairness axioms such as Pareto optimality<sup>7</sup> [Nash 1950]) to compute bandwidth allocations for elastic services and pricing in broadband networks. This framework provides rate allocations for users that are not only Pareto optimal from the viewpoint of the whole system but also are consistent with game-theoretic fairness axioms.

Other researchers have used game-theoretic formulations to investigate the effect of pricing on user adoption and fair use. Pricing models to induce participation and collaboration in a public wireless mesh network were studied by Lam et al. [2007]. Shen and Basar [2007] investigated optimal nonlinear pricing policy design as a means of controlling network usage and generating profits for a monopolistic service provider. Dynamic game models have also been used to determine WiFi access point pricing by Musacchio and Walrand [2006], which is relevant in the context of congestion management through WiFi offloading.

Jiang et al. [2008] introduced a model to study the role of time preferences in network pricing. In their model, each user chooses his or her access time based on his or her preference, the congestion level, and the price charged. The authors show that maximization of both the social welfare and the revenue of a service provider is feasible if the provider can differentiate its prices over different users and times. However, if the prices can only be differentiated over the access times and not across users due to insufficient information, the resulting social welfare can be much less than the optimum, especially in the presence of many low-utility users. Caron and Kesidis [2010] took a similar approach for the electricity market; in their model, users schedule their energy usage in a cooperative game so as to minimize the total load on the network.

Despite these theoretical works, game-theoretic models have found little traction among real network operators so far, perhaps due to the stylized nature of the theoretical models and the challenges in estimating user utility and system parameters in the real world.

---

<sup>7</sup>A Pareto-optimal allocation of resources is one for which increasing any one user's allocation will decrease another user's allocation.

Table I. List of Key Example Pricing Plans

| Type                         | Pricing Practice  |  | Example Pricing Plan (see the article for details) |   |
|------------------------------|-------------------|--|--|---|
|                              | Category          | Description                              | Network  | Country                                 |
| Static                       | Fixed Flat Rate   | Monthly fee; unlimited                   | Both   | Vanishing worldwide                     |
|                              |                   | Monthly; flat to a cap, then usage based | Wired  | U.S. (AT&T, Verizon)                    |
|                              |                   | Monthly; flat to a cap, then throttle    | Wired/wireless                                     | Spain (Orange)<br>U.S. (Comcast)        |
|                              |                   | Monthly; shared                          | Wired/wireless                                     | Canada (Rogers)<br>U.S. (AT&T, Verizon) |
|                              | Usage Based       | Hourly rate                              | Wireless   | Egypt (Mobinil)                         |
|                              |                   | Cap, then metered                        | Wired/wireless                                     | Worldwide (e.g., U.S., U.K.)            |
|                              | Priority Pricing  | Priority pass (for dongle users)         | Wireless   | Singapore (SingTel)                     |
|                              | Time of Day       | Daytime & nighttime rates                | Wireless   | India (BSNL)                            |
|                              |                   | Users choose happy hours                 | Wireless   | U.K. (Orange)                           |
|                              | Cumulus Pricing   | Usage-based contract negotiation         | Wireless   | U.K. (Vodafone)                         |
|                              | App-Based Pricing | Free access to select apps; bundling     | Wireless   | U.K. (Orange)<br>Denmark (TDC)          |
|                              | Dynamic           | Congestion Based                         | Hourly price changes                               | Wireless (voice calls)                  |
| Location and cell-load based |                   |  | Wireless (voice calls)                             | India (Uninor)                          |
| Time- and usage based        |                   |  | Wireless (data)                                    | U.S. (pilot trial)                      |

The different real-world examples of various static and dynamic pricing plans for the wired and wireless services discussed in this article are summarized in Table I.

## 7. EMERGING TECHNOLOGY TRENDS

The recent exponential growth in data demand has catalyzed major changes within broadband pricing research and practice. In this section, we discuss the new directions that are emerging from these recent changes.

### 7.1. Satellite Broadband

This survey focuses mainly on wired and wireless pricing, as these are the mediums most impacted by users' increasing demands for data. Recently, however, many companies have begun to offer satellite broadband as an alternative to wired or wireless Internet access. Although satellite broadband solutions have existed for more than 20 years, satellite Internet access has only recently become a popular offering for end users. Satellite is becoming especially prevalent as a way to reach users in rural and sparsely populated regions, such as in Africa, where installing more traditional wired or wireless infrastructures is not cost-effective [Svitak 2012].

Table II. Satellite Plan Features and Congestion Control Mechanisms

| Feature                    | Options  |
|----------------------------|--|
| Data caps (limits)         | Monthly [Muller 2012]<br>Daily   |
| Uplink and downlink speeds | Varies by company [D’Orazio 2012; IPNet JSC 2012]  |
| Overage penalty            | Throttling to lower speeds [D’Orazio 2012]<br>Usage-based charges [IPNet JSC 2012; Rural Broadband 2012] |

Since satellite broadband has not yet experienced serious congestion problems, its pricing plans remain fairly simple, although the growth in demand for high-bandwidth services may open up possibilities for new pricing schemes in the future. Technological capabilities that allow for fine granular capacity provisioning and dynamic channel allocation will also facilitate the adoption of more innovative pricing practices. Sun and Modiano [2006] proposed a pricing-based solution to the problem of channel allocation in slotted Aloha-based satellite networks. Today, satellite broadband consumers mostly pay in proportion to the uplink and downlink data rates by subscribing to a chosen tier of available data plans. Each tier has specified maximum upload and download speeds and a monthly cap on bandwidth. Beyond this cap, a user’s speed is either throttled down or charged by overage, depending on the plan. Table II gives an overview of key features of different plans, which vary by the granularity of data caps, the offered uplink and downlink speeds, and the type of overage penalties.

## 7.2. Pricing of Heterogeneous Networks

As demand for mobile data grows, ISPs are not only introducing new pricing plans but also are encouraging users to offload their traffic onto other wireless networks. For instance, AT&T deployed WiFi hotspots in New York and San Francisco after their cellular networks proved unable to handle the demand for data in such congested urban areas [Ricker 2010]. Also in the United States, Verizon stated as early as 2012 that it expected to run out of LTE spectrum in the next two years in some markets; the company is encouraging consumer adoption of femtocells to supplement this limited capacity [Electronista 2012]. ISPs’ key reasons for rolling out small cells are that they help increase capacity supply while keeping costs low and rely on end-user backhauling for data transfers. Moreover, higher band spectrum is now becoming available, which is suitable for small cells, in particular femtocells. Although small cells are cheap to deploy and work as plug-and-play devices, they can often lead to poor performance due to interference on the shared spectrum, which poses considerable technical challenge. With the increasing adoption of small cells and their coexistence with other wireless networks (e.g., 3G, 4G), ISPs will soon need to consider the pricing of heterogeneous networks with devices being able to connect to either of the available network technologies.

As with all pricing, heterogeneous network pricing can be used to steer users to an optimal allocation of resources that maximizes the efficiency of the networks considered. For instance, Neely et al. [2005] examined flow control in a heterogeneous network setting with wired and wireless services. They employed methods from stochastic optimal control to optimize user utility subject to the bandwidth capacity constraints, an approach that may be interpreted as a type of dynamic market pricing. Chan et al. [2005] assumed that users have access to multiple networks, possibly more than two, and used dynamic pricing to steer users to the “optimal” network so as to maximize overall user utility.

Other, more recent works on heterogeneous networks have emphasized the profit and revenue of ISPs, rather than employing prices solely to optimize resource management. These works often begin by modeling user adoption of heterogeneous technologies as a function of the access prices; users are assumed to adopt different technologies



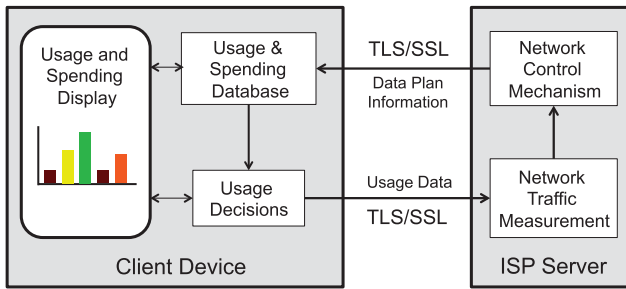


Fig. 7. Pricing architecture with client- and server-side components.

so as to maximize their utility, which in turn depends on the total number of users adopting each technology (e.g., due to positive externalities of connectivity, or negative externalities from network congestion). Shetty et al. [2009] used such a model to study revenue maximization for femtocells, whereas Ren et al. [2013] examined the selection of femtocell spectrum-sharing schemes and the corresponding optimal prices. Joe-Wong et al. [2013] took a more generic approach, studying the adoption of generic “base” and “supplemental” technologies and the ISP’s profit-maximizing prices when offloading benefits and network deployment costs are considered. A game-theoretic formulation that does not include adoption modeling was proposed by Niyato and Hossain [2008], in which two companies competing for end users were assumed to offer WiFi and WIMAX connectivity.

### 7.3. New Pricing Architectures

Academics have often argued that architectural issues in Internet pricing are more important than the form or basis of the actual prices charged: for instance, Shenker et al. [1996] opined that the flat- versus usage-based pricing debate oversimplifies the issue, as in reality there is a continuum between the two. The authors made the case that research into Internet pricing has overly focused on optimality, and in particular on matching prices to the marginal congestion costs. Instead, they advocated an architectural focus—that is, designing the network architecture to facilitate various pricing plans, such as allowing receivers rather than senders to be charged for usage. Architecture and systems enabling such pricing innovations for broadband data are arguably becoming an active area of research and are central to shaping the future research agenda. Understanding these interactions between economics and technology in order to create the overall design of future networks has been identified as a priority in the NSF’s Future Internet Design initiative [Fisher 2007].

As new pricing plans are introduced, ISPs must confront the challenges of ensuring that they are acceptable to users, while also maintaining system scalability, privacy, and security. One way to do this, such as that proposed by MacKie-Mason et al. [1997], is to add a client-side component to the current pricing architecture, as shown in Figure 7. The client-side component includes modules to show users their data consumption and spending over the month, thus allowing users to monitor and control their spending, even if they are unfamiliar with their data plan specifics. Secure connections with the ISP server ensure that user privacy is not compromised. In fact, data-monitoring apps that sit exclusively on the client side are already being offered on both wireless and wired platforms [Princeton EDGE Lab 2012; Locktime Software 2012] to help users manage monthly data caps. On wired desktop platforms, users can even control their data usage by controlling the bandwidth rates of different applications. Research on in-home Internet usage has showed that such applications enhance user awareness of

data usage and that there is a consumer demand for tools that help them understand their data usage [Chetty et al. 2010, 2012]. However, the interfaces of these applications for both wired and wireless networks must be carefully designed in order to facilitate user engagement [Chetty et al. 2011; Sen et al. 2013].

Client-side usage-monitoring applications can go far beyond monitoring usage, especially if integrated with an ISP pricing plan. For instance, the client-side architecture for dynamic pricing can incorporate a module to display current and, if known, future prices as well as usage and spending. In fact, if future prices are known, such as with day-ahead dynamic pricing, one can even go a step further and incorporate a module that automatically schedules some delay-tolerant applications to less expensive times of the day [Ha et al. 2012b]. Such an “autopilot mode” for usage decisions incorporates a form of bandwidth control into wireless devices, and integrating it with the prices offered can help make dynamic pricing viable for ISPs.

#### 7.4. Emerging Pricing Plans

A new pricing option that U.S. ISPs are offering is a “shared data plan,” which allows users to share their data cap across multiple devices, at a premium for each additional device [Buckley 2012; Galbraith 2011]. Rogers in Canada has been offering such plans on a promotional basis since 2011 [Electronista 2011]. Orange of France Telecom has also started such a plan in Austria that allows iPad owners to share one allotment of data with a phone, and 38% of iPad owners on its network now subscribe to this plan [Fried 2011]. Sharing a data cap across multiple devices may induce more efficient usage of the quota per consumer. However, an open research question is to understand shared data plans’ impact on usage and quota dynamics, with field trials to understand the trade-offs between fairly sharing the caps across devices.

In dealing with the growing problem of bandwidth demand, researchers from Alcatel-Lucent have recently proposed static time-dependent discounts for incentivizing off-peak capacity usage [El-Sayed et al. 2011]. They also presented policy and QoS mechanisms to implement traffic shaping with off-peak usage incentives to users. Another interesting trend emerging in several parts of the world, particularly in Asia and Africa, is that of a larger-scale adoption of dynamic pricing for voice calls [The Economist 2009]. As the demand for data, and in particular mobile broadband, grows, there will be a need for new architectures and systems, as reported by Ha et al. [2012] and Joe-Wong et al. [2011] in the context of pilot-trial results for a dynamic time-dependent usage-based pricing system called TUBE (Time-dependent Usage-based Broadband price Engineering) [DataMi 2012; Ha et al. 2012b].

Differentiated pricing for broadband [Ericsson 2011] (e.g., application-based pricing of data) is also a trend that is likely to receive more attention from industry. As noted previously, ISPs are exploring new ways to attract customers by “zero-rating” or bundling certain applications. A related direction for future research is that of data plans with sponsored content in which users’ access fees are subsidized by the content providers or advertisers. By offering such plans, content providers can incentivize users to consume their content, and thus increase their advertising revenue. Sponsored content pricing can be modeled as a two-sided market; using game theory, Andrews et al. [2013] derived the optimal pricing policies for an ISP. However, such a pricing scheme could be viewed as nonneutral, which may hinder its adoption. To avoid this issue, Sheng et al. [2013] suggested a platform that allows content providers to subsidize a certain dollar amount of users’ traffic, which need not correspond to traffic from that content provider. In exchange, users agree to view advertisements, which can help generate content provider revenue.

A summary of these potential future trends is given in Table III.

Table III. Summary of Some Emerging Trends

| Trend                             | Description   | Reference   |
|-----------------------------------|---|---|
| Flat-rate & usage-based pricing   | Coexistence of both types of plans.<br>Emergence of shared data plans.<br>Focus on architectural issues in Internet pricing is needed.                                      | [Buckley 2012; Fried 2011]<br>[Shenker et al. 1996]                 |
| Time-dependent pricing            | Traffic shaping with off-peak usage incentives.<br>Static discounts (two-period ToD).   | [El-Sayed et al. 2011]  |
| Dynamic pricing                   | Practiced for voice calls (e.g., location based).<br>Adopted in Africa & India.<br>Needs new theory for dynamic data pricing.<br>New architecture & systems implementation. | [The Economist 2009]<br>[Joe-Wong et al. 2011]<br>[Ha et al. 2012b] |
| Differentiated, app-based pricing | Tiered data plans with various QoS.<br>App-specific “zero-rating,” app bundles.   | [Ericsson 2011]<br>[Goldstein 2012]                                 |

## 8. CONCLUSIONS

In this work, we draw attention to the growing problem of network congestion and highlight some of the recent steps taken by ISPs to mitigate its effects. The projected growth in demand for data, especially from mobile data and video traffic, is far more than what can be supported even by the latest technological advances (e.g., 4G/LTE and WiFi offloading) due to expensive backhauling and increasing wired network congestion. Consequently, ISPs have been aggressively using pricing as a congestion control tool.

The basic idea of congestion pricing has been known in the networking community for several decades, but only now have conditions demanded that it be put into practice. In this survey, we first outline the various threats to the sustainability and economic viability of the Internet ecosystem from the perspectives of ISPs, consumers, and content providers. We then highlight the role that pricing can play in alleviating network congestion and in balancing the functionality goals of data networks, followed by a detailed discussion of the technological, socioeconomic, and regulatory challenges to pricing innovation and related open problems. We review many known data network pricing proposals, both static and dynamic, and discuss the extent to which some of these have been adopted by ISPs, drawing parallels to existing pricing practices in electricity markets and road networks. We also discuss the predominant trends in access pricing and the need for new architecture, field trials, and interdisciplinary research for realizing more innovative dynamic pricing solutions such as day-ahead time- and usage-dependent pricing and app-based pricing.

The material presented in this article seeks to inform networking researchers about the existing works on access pricing, the ongoing developments in pricing plans, and the research challenges that need to be accounted for in shaping a new research agenda on smart data pricing.

## ACKNOWLEDGMENTS

This work has benefited immensely from our discussions with several colleagues and collaborators. We would like to specially thank (in no particular order) Victor Glass, Prashanth Hande, Krishan Sabnani, Raj Savor, Steve Sposato, Rob Calderbank, Danny Tsang, Sundeep Rangan, Yuan Wu, Junshan Zhang, T. Russell Hsing, Keith Cambron, Andrew Odlyzko, Roch Gu erin, Augustin Chaintreau, and Jennifer Rexford for their valuable comments, viewpoints, and information that have helped to shape this article. Additionally, we acknowledge the support of our industrial collaborators, in particular, Reliance Communications of India, MTA of Alaska, SES, AT&T, Qualcomm, Comcast, and the National Exchange Carrier Association (NECA). We also acknowledge the generous feedback received from the participants of the first Workshop on Smart Data Pricing [SDP Forum 2012], held in Princeton, NJ, on July 30–31, 2012.

## REFERENCES

- AL-MANTHARI, B., NASSER, N., AND HASSANEIN, H. 2011. Congestion pricing in wireless cellular networks. *IEEE Commun. Surv. Tutor.* 13, 3, 358–371.
- ALLOT MOBILETRENDS. 2010. Global mobile broadband traffic report.
- ANANIA, L. AND SOLOMON, R. J. 1997. Flat—the minimalist price. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 91–118.
- ANDREWS, M., ÖZEN, U., REIMAN, M. I., AND WANG, Q. 2013. Economic models of sponsored content in wireless networks with uncertain demand. In *Proceedings of the 2nd IEEE International Workshop on Smart Data Pricing (SDP'13)*. IEEE, 345–350.
- APPLE. 2012. Apple—the new iPad. <http://www.apple.com/ipad/compare/>.
- BORENSTEIN, S., JASKE, M., AND ROSENFELD, A. 2002. Dynamic pricing, advanced metering, and demand response in electricity markets. Working paper CSEM-105, Center for the Study of Energy Markets, Berkeley, CA.
- BREKER, L. P. 1996. A survey of network pricing schemes. In *Proceedings of the 8th Symposium on Computer Science*.
- BROWNLIE, N. 1997. Internet pricing in practice. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 77–90.
- BUBLEY, D. 2011. Belgian MNO tries app-specific zero rating—bad idea in my view. Disruptive Wireless. November 1, <http://tinyurl.com/6urunl3>.
- BUCKLEY, S. 2012. Verizon shared data plans show up in employee training materials, still on track? Engadget. January 30.
- CAREW, S. 2012. Users complain, AT&T blames data tsunami. Reuters Media File.
- CARON, S. AND KESIDIS, G. 2010. Incentive-based energy consumption scheduling algorithms for the smart grid. In *Proceedings of the 1st IEEE International Conference on Smart Grid Communications*. IEEE, Los Alamitos, CA, 391–396.
- CHAN, H., FAN, P., AND CAO, Z. 2005. A utility-based network selection scheme for multiple services in heterogeneous networks. In *Proceedings of the International Conference on Wireless Networks, Communications and Mobile Computing: Volume 2*. IEEE, 1175–1180.
- CHANG, X. AND PETR, D. 2001. A survey of pricing for integrated service networks. *Comput. Commun.* 24, 18, 1808–1818.
- CHARLES RIVER ASSOCIATES. 2005. Impact evaluation of the California statewide pricing pilot. Tech. rep., Charles River Associates.
- CHEN, B. X. 2012. Shared mobile data plans: Who benefits? *New York Times*. July 19, Bits Blog.
- CHEN, J., GHOSH, A., MAGUTT, J., AND CHIANG, M. 2012. QAVA: Quota aware video adaptation. In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies (CoNEXT'12)*. ACM, 121–132.
- CHEN, Y.-F. R. AND JANA, R. 2013. SpeedGate: A smart data pricing testbed based on speed tiers. In *Proceedings of the 2nd Smart Data Pricing Workshop*. IEEE, 3195–3200.
- CHETTY, M., BANKS, R., BRUSH, A., DONNER, J., AND GRINTER, R. 2012. You're capped: Understanding the effects of bandwidth caps on broadband use in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'12)*. ACM, 3021–3030.
- CHETTY, M., BANKS, R., HARPER, R., REGAN, T., SELLEN, A., GKANTSIDIS, C., KARAGIANNIS, T., AND KEY, P. 2010. Who's hogging the bandwidth: The consequences of revealing the invisible in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'10)*. ACM, 659–668.
- CHETTY, M., HASLEM, D., BAIRD, A., OFOHA, U., SUMNER, B., AND GRINTER, R. 2011. Why is my Internet slow?: Making network speeds visible. In *Proceedings of SIGCHI Conference on Human Factors in Computing Systems (CHI'11)*. ACM, 1889–1898.
- CISCO SYSTEMS. 2012a. Cisco visual networking index: Forecast and methodology, 2011–2016. May 30, <http://tinyurl.com/VNI2011>.
- CISCO SYSTEMS. 2012b. Cisco visual networking index: Global mobile data traffic forecast update, 2011–2016. February 14, <http://tinyurl.com/VNI2012-mobile>.
- CLARK, D. D. 1997. Internet cost allocation and pricing. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 215–252.
- COCCHI, R., ESTRIN, D., SHENKER, S., AND ZHANG, L. 1991. A study of priority pricing in multiple service class networks. *ACM SIGCOMM Comput. Commun. Rev.* 21, 4, 123–130.
- COCCHI, R., SHENKER, S., ESTRIN, D., AND ZHANG, L. 1993. Pricing in computer networks: Motivation, formulation, and example. *IEEE/ACM Trans. Networking.* 1, 614–627.

- COMCAST. 2012. What XFINITY Internet data usage plans will Comcast be launching? <http://customer.comcast.com/help-and-support/internet/data-usage-what-are-the-different-plans-launching>.
- COURCOUBETIS, C., STAMOULIS, G. D., MANOLAKIS, C., AND KELLY, F. P. 1998. An intelligent agent for optimizing QoS-for-money in priced ABR connections. Tech. rep., Institute of Computer Science–Foundation for Research and Technology Hellas and Statistical Laboratory, University of Cambridge. Preprint.
- DASH, D., KANTERE, V., AND AILAMAKI, A. 2009. An economic model for self-tuned cloud caching. In *Proceedings of the IEEE 25th International Conference on Data Engineering (ICDE'09)*. IEEE, Los Alamitos, CA, 1687–1693.
- DATA MI. 2012. DataMi Website. <http://scenic.princeton.edu/datami/>.
- DELGROSSI, L. AND FERRARI, D. 1999. Charging schemes for reservation-based networks. *Telecommun. Syst.* 11, 1, 127–137.
- D'ORAZIO, D. 2012. ViaSat prices satellite broadband plans, calls them Exede. TheVerge. January 10, <http://www.theverge.com/2012/1/10/2694708/viasat-exede-pricing-availability-satellite-broadband-internet>.
- DU, P. AND LU, N. 2011. Appliance commitment for household load scheduling. *IEEE Trans. Smart Grid* 2, 2, 411–419.
- DYABERI, J. M., PARSONS, B., PAI, V. S., KANNAN, K., CHEN, Y.-F. R., JANA, R., STERN, D., AND VARSHAVSKY, A. 2012. Managing cellular congestion using incentives. *IEEE Commun. Mag.* 50, 11, 100–107.
- EL-SAYED, M., MUKHOPADHYAY, A., URRUTIA-VALDÉS, C., AND ZHAO, Z. J. 2011. Mobile data explosion: Monetizing the opportunity through dynamic policies and QoS pipes. *Bell Labs Techn. J.* 16, 2, 79–100.
- ELECTRONISTA. 2011. Rogers intros data sharing for 6GB plans with hefty premium. February 8, <http://www.electronista.com/articles/11/02/08/rogers.6gb.data.plan.gets.expensive.sharing/>.
- ELECTRONISTA. 2012. Verizon preps LTE femtocells, says it hits limits in 2013. March 5, <http://www.electronista.com/articles/12/03/05/demand.for.mobile.data.will.outpace.build.out/>.
- ERICSSON. 2010. Japan's Softbank Mobile chooses Ericsson mobile broadband solution. August 2, press release.
- ERICSSON. 2011. Differentiated mobile broadband. White paper, [http://www.ericsson.com/res/docs/whitepapers/differentiated\\_mobile\\_broadband.pdf](http://www.ericsson.com/res/docs/whitepapers/differentiated_mobile_broadband.pdf).
- EZZIANE, Z. 2005. Charging and pricing challenges for 3G systems. *IEEE Commun. Surv. Tutor.* 7, 4, 58–68.
- FALKNER, M., DEVETSIKIOTIS, M., AND LAMBADARIS, I. 2000. An overview of pricing concepts for broadband IP networks. *IEEE Commun. Surv. Tutor.* 3, 2, 2–13.
- FARUQUI, A. 2010. The ethics of dynamic pricing. *Electricity J.* 23, 6, 13–28.
- FARUQUI, A., HLEDIK, R., AND SERGICI, S. 2009. Piloting the smart grid. *Electricity J.* 22, 7, 55–69.
- FARUQUI, A. AND WOOD, L. 2008. Quantifying the benefits of dynamic pricing in the mass market. Tech. rep., Edison Electric Institute.
- FISHER, D. 2007. US National Science Foundation and the future Internet design. *ACM SIGCOMM Comput. Comm. Rev.* 37, 3, 85–87.
- FITCHARD, K. 2012. New Netflix iOS app capitulates to bandwidth caps. GigaOm. May 31, <http://gigaom.com/mobile/new-netflix-ios-app-capitulates-to-bandwidth-caps/>.
- FRAKES, D. 2010. AT&T announces tethering details and new plans for iPhone, iPad. InfoWorld.
- FRIED, I. 2011. The data plan of the future is available now, at least in Europe. All Things D. <http://tinyurl.com/FT-shared>.
- GALBRAITH, C. 2011. Shared data plans set to make global impact. Billing World. October 24.
- GANESH, A., LAEVEENS, K., AND STEINBERG, R. 2001. Congestion pricing and user adaptation. In *Proceedings of the 20th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'01): Volume 2*. IEEE, 959–965.
- GENACHOWSKI, J. 2010. New rules for an open Internet. US Federal Communications Commission. December 21, <http://www.fcc.gov/blog/new-rules-open-internet>.
- GIZELIS, C. AND VERGADOS, D. 2011. A survey of pricing schemes in wireless networks. *IEEE Commun. Surv. Tutor.* 13, 1, 126–145.
- GLASS, V., PRINZIVALLI, J., AND STEFANOVA, S. 2012. Persistence of the middle mile problem for rural exchanges local carriers. Smart Data Pricing Workshop Talk. <http://scenic.princeton.edu/SDP2012/Talks-VictorGlass.pdf>.
- GOLDSTEIN, P. 2012. AT&T's Stephenson: Content providers are asking for 'toll free' data plans. FierceWireless. <http://www.fiercewireless.com/story/atts-stephenson-content-providers-are-as-king-toll-free-data-plans/2012-06-01#ixzz2M466tpzz>.
- GOMEZ-IBANEZ, J. A. AND SMALL, K. A. 1994. *Road Pricing for Congestion Management: A Survey of International Practice*. Transportation Research Board, Washington, DC. National Cooperative Highway Research Program.



- GUPTA, A., STAHL, D., AND WHINSTON, A. 1997. Priority pricing of integrated services networks. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 323–352.
- HA, S., JOE-WONG, C., SEN, S. AND CHIANG, M. 2012a. Pricing by timing: Innovating broadband data plans. In *Proceedings of the SPIE OPTO Broadband Access Communication Technologies VI Conference (SPIE 8282)*. Paper 8282-12.
- HA, S., SEN, S., JOE-WONG, C., IM, Y., AND CHIANG, M. 2012b. TUBE: Time-dependent pricing for mobile data. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM'12)*. ACM, New York, NY, 247–258.
- HANDE, P., CHIANG, M., CALDERBANK, R., AND ZHANG, J. 2010. Pricing under constraints in access networks: Revenue maximization and congestion management. In *Proceedings of the IEEE INFOCOM*. IEEE, 1–9.
- HARSHMAN, B. 2001. Urban road pricing acceptance. In *Proceedings of the IMPRINT-EUROPE Seminar*.
- HAUSMAN, J., KINNUCAN, M., AND MCFADDEN, D. 1979. A two-level electricity demand model: Evaluation of the Connecticut time-of-day pricing test. *J. Econometr.* 10, 3, 263–289.
- HAYEL, Y. AND TUFFIN, B. 2005. A mathematical analysis of the cumulus pricing scheme. *Comput. Netw.* 47, 907–921.
- HAYER, J. 1993. Transportation auction: A new service concept. Master's thesis, University of Alberta. TR-93-05.
- HENDERSHOTT, T. 2006. *Economics and Information Systems, Volume 1*. Elsevier, Amsterdam, Chapter 2, 4–9.
- HERTER, K. 2007. Residential implementation of critical-peak pricing of electricity. *Energy Policy* 35, 4, 2121–2130.
- HIGGINBOTHAM, S. 2010. Mobile operators want to charge based on time and apps. GigaOm. December 14.
- HOLGUIN-VERAS, J., CETIN, M., AND XIA, S. 2006. A comparative analysis of US toll policy. *Transport. Res. A-Pol.* 40, 10, 852–871.
- HOSANAGAR, K., KRISHNAN, R., CHUANG, J., AND CHOUDHARY, V. 2005. Pricing and resource allocation in caching services with multiple levels of quality of service. *Manage. Sci.* 51, 12, 1844–1859.
- IBM. 2007. Ontario energy board smart price pilot final report. Tech. rep., Ontario Energy Board. IBM Global Business Services and eMeter Strategic Consulting.
- IPNET JSC. 2012. Price and rates. <http://www.zaoinet.ru/en/price/connectionservice/>.
- JIANG, L., PAREKH, S., AND WALRAND, J. 2008. Time-dependent network pricing and bandwidth trading. In *Proceedings of the IEEE Network Operations and Management Symposium Workshop*. IEEE, 193–200.
- JOE-WONG, C., HA, S., AND CHIANG, M. 2011. Time-dependent broadband pricing: Feasibility and benefits. In *Proceedings of the 31st International Conference on Distributed Computing Systems (ICDCS'11)*. IEEE, 288–298.
- JOE-WONG, C., SEN, S., AND HA, S. 2013. Offering supplementary wireless network technologies: Adoption behavior and offloading benefits. In *Proceedings of the IEEE INFOCOM*. <http://arxiv.org/abs/1209.5004>.
- JOE-WONG, C., SEN, S., HA, S., AND CHIANG, M. 2012a. Optimized day-ahead pricing for smart grids with device-specific scheduling flexibility. *IEEE J. Sel. Areas Commun.* 30, 6, 1075–1085.
- JOE-WONG, C., SEN, S., LAN, T., AND CHIANG, M. 2012b. Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework. In *Proceedings of the IEEE INFOCOM*. IEEE, 1206–1214.
- KANG, C. 2010. AT&T wireless scraps flat-rate Internet plan. Washington Post. June 2.
- KEALL, C. 2011. TelstraClear calls limitless weekend 'very successful,' InternetNZ 'disastrous.' National Business Review. December 5.
- KELLY, F. 1994. On tariffs, policing, and admissions control for multiservice networks. *Oper. Res. Lett.* 15, 1–9.
- KELLY, F., MAULLOO, A. K., AND TAN, D. H. K. 1998. Rate control for communication networks: Shadow prices, proportional fairness, and stability. *J. Oper. Res. Soc.* 49, 237–252.
- KEY, P. 2010. Comcast, Level 3, Netflix, the FCC: Busy week for neutrality debate. Philadelphia Business Journal. December 1.
- KLOPFENSTEIN, B. C. 2009. Internet economics: An annotated bibliography. *J. Media Econ.* 11, 1, 33–48.
- KWANG, K. 2011. SingTel unveils priority mobile broadband access. June 13, <http://www.zdnetasia.com/singtel-unveils-priority-mobile-broadband-access-62300719.htm>.
- LAM, K. R., CHIU, D. M., AND LUI, J. C. S. 2007. On the access pricing and network scaling issues of wireless mesh networks. *IEEE Trans. Comput.* 56, 140–146.
- LAOUTARIS, N., SIRIVIANOS, M., YANG, X., AND RODRIGUEZ, P. 2011. Inter-datacenter bulk transfers with netstitcher. *ACM SIGCOMM Comput. Comm. Rev.* 41, 4, 74–85.
- LEE, D., MO, J., WALRAND, J., AND PARK, J. 2011. A token pricing scheme for Internet services. In *Proceedings of the 7th International Conference on Internet Charging and QoS Technologies: Economics of the Converged, Internet-Based Networks (ICQT'11)*. 26–37.

- LEWIS, P. H. 1996. An “all you can eat” price is clogging Internet access. *New York Times*. December 17, p. A1.
- LI, N., CHEN, L., AND LOW, S. 2011. Optimal demand response based on utility maximization in power networks. In *Proceedings of the IEEE Power and Energy Society General Meeting*. IEEE, Los Alamitos, CA, 1–8.
- LI, S., HUANG, J., AND LI, S. Y. R. 2009. Revenue maximization for communication networks with usage-based pricing. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM'09)*. IEEE, 1–6.
- LOCKTIME SOFTWARE. 2012. NetLimiter Website. <http://www.netlimiter.com/>.
- LOISEAU, P., SCHWARTZ, G., MUSACCHIO, J., AND AMIN, S. 2011. Incentive schemes for Internet congestion management: Raffles versus time-of-day pricing. In *Proceedings of the 49th Annual Allerton Conference on Communication, Control, and Computing*. IEEE, 103–110.
- MACKIE-MASON, J. AND VARIAN, H. 1995. Pricing the Internet. In B. Kahin and J. Keller, Eds., *Public Access to the Internet*. Prentice Hall, Englewood Cliffs, NJ.
- MACKIE-MASON, J. K., MURPHY, L., AND MURPHY, J. 1997. Responsive pricing in the Internet. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 279–303.
- MARBACH, P. 2004. Analysis of a static pricing scheme for priority services. *IEEE/ACM Trans. Netw.* 12, 312–325.
- MARCON, M., SANTOS, N., GUMMADI, K. P., LAOUTARIS, N., RODRIGUEZ, P., AND VAHDAT, A. 2012. NetEx: Cost-effective bulk data transfers for cloud computing. Tech. rep., Max Planck Institute for Software Systems.
- MARSHALL, J. 1996. Economics, not engineering, will unclog the Internet. *San Francisco Chronicle*. November 4.
- MATSUKAWA, I. 2001. Household response to optional peak-load pricing of electricity. *J. Regul. Econ.* 20, 3, 249–267.
- MAY, A. 1992. Road pricing: An international perspective. *Transportation* 19, 4, 313–333.
- MAY, A. AND MILNE, D. 2000. Effects of alternative road pricing systems on network performance. *Transport. Res. A-Pol.* 34, 6, 407–436.
- McKINNON, A. 2006. A review of European truck tolling schemes and assessment of their possible impact on logistics systems. *Int. J. Logist. Res. Appl.* 9, 3, 191–205.
- McKNIGHT, L. W. AND BAILEY, J. P. 1997. An introduction to Internet economics. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA, 91–118.
- MOHSENIAN-RAD, A. H. AND LEON-GARCIA, A. 2010. Optimal residential load control with price prediction in real-time electricity pricing environments. *IEEE Trans. Smart Grid* 1, 2, 120–133.
- MOHSENIAN-RAD, A. H., WONG, V. W. S., JATSKEVICH, J., AND SCHOBER, R. 2010. Optimal and autonomous incentive-based energy consumption scheduling algorithm for smart grid. In *Proceedings of the IEEE Conference on Innovative Smart Grid Technologies*. IEEE, Los Alamitos, CA, 1–6.
- MORGAN, M. 2011. Pricing schemes key in LTE future. *Telecomasia.net*. September 12, <http://www.telecomasia.net/content/pricing-schemes-key-lte-future>.
- MULLER, R. 2012. Satellite broadband pricing comparison. January 25, <http://mybroadband.co.za/news/broadband/>.
- MURPHY, J. AND MURPHY, L. 1994. Bandwidth allocation by pricing in ATM networks. *IFIP Trans. C. C-24*, 333–351.
- MURPHY, J., MURPHY, L., AND POSNER, E. C. 1994. Distributed pricing for embedded ATM networks. In *Proceedings of the International Teletraffic Conference*. Elsevier, 1053–1063.
- MUSACCHIO, J. AND WALRAND, J. 2006. WiFi access point pricing as a dynamic game. *IEEE/ACM Trans. Netw.* 14, 289–301.
- NASH, J. F. 1950. The bargaining problem. *Econometrica* 18, 2, 155–162.
- NEELY, M., MODIANO, E., AND LI, C. 2005. Fairness and optimal stochastic control for heterogeneous networks. In *Proceedings of the IEEE INFOCOM: Volume 3*. IEEE, Los Alamitos, CA, 1723–1734.
- NEWMAN, J. 2011. Netflix has bandwidth cap sufferers covered. *PC World*. June 23, [http://www.pcworld.com/article/230982/Netflix\\_Has\\_Bandwidth\\_Cap\\_Sufferers\\_Covered.html](http://www.pcworld.com/article/230982/Netflix_Has_Bandwidth_Cap_Sufferers_Covered.html).
- NetEcon. 2013. Workshop on the economics of networks, systems and computation. <http://netecon.seas.harvard.edu/>.
- NICHOLSON, W. AND SNYDER, C. 2008. *Microeconomic Theory* (10th ed.). SouthWestern CENGAGE Learning, Mason, OH.
- NIYATO, D. AND HOSSAIN, E. 2008. Competitive pricing in heterogeneous wireless access networks: Issues and approaches. *IEEE Netw.* 22, 6, 4–11.
- ODLYZKO, A. 1999. Paris metro pricing for the Internet. In *Proceedings of the 1st ACM Conference on Electronic Commerce*. ACM, New York, 140–147.
- ODLYZKO, A. 2001. Internet pricing and the history of communications. *Comput. Netw.* 36, 5, 493–517.

- ODLYZKO, A., ARNAUD, B. S., STALLMAN, E., AND WEINBERG, M. 2012. Know your limits: Considering the role of data caps and usage based billing in Internet access service. Public Knowledge.
- OWOSEJE, T. 2011. Orange refreshes Panther tariff with Swapables benefits. Mobile Magazine. September 7, <http://tinyurl.com/Swapables>.
- PARKER, A. 2010. O2 scraps unlimited data for smartphones. Financial Times.
- PARRIS, C. AND FERRARI, D. 1992. A resource based pricing policy for real-time channels in a packet-switching network. Tech. rep. TR-92-016. Tenet Group, ICSI, University of California, Berkeley, CA.
- PARRIS, C., KESHAV, S., AND FERRARI, D. 1992. A framework for the study of pricing in integrated networks. Tech. rep. TR-92-016. Tenet Group, ICSI, University of California, Berkeley, CA.
- PASCHALIDIS, I. C. AND TSITSIKILIS, J. N. 1998. Congestion-dependent pricing of network services. *IEEE/ACM Trans. Netw.* 8, 171–184.
- POOLE, R. AND ORSKI, C. K. 2003. HOT networks: A new plan for congestion relief and better transit. Tech. rep., Reason Foundation. February.
- PRINCETON EDGE LAB. 2012. DataWiz Website. <http://scenic.princeton.edu/datawiz/>.
- REICHL, P. 2010. From charging for quality of service to charging for quality of experience. *Ann. Telecommun.* 65, 3, 189–199.
- REN, S., PARK, J., AND VAN DER SCHAAR, M. 2013. Entry and spectrum sharing scheme selection in femtocell communications markets. *IEEE/ACM Trans. Netw.* 21, 1, 218–232.
- RICHARDS, M., GILLIAM, C., AND LARKINSON, J. 1996. The London congestion charging research programme. 1. The programme in overview. *Traffic Eng. Control* 37, 2, 66–71.
- RICKER, T. 2010. AT&T making tourists even more annoying with free Times Square WiFi. Engadget. <http://tinyurl.com/29quo58>.
- RIGNEY, C., WILLENS, S., RUBENS, A., AND SIMPSON, W. 2000. Remote Authentication Dial In User Service (RADIUS). RFC 2865 (Draft Standard). Updated by RFCs 2868, 3575, 5080.
- ROOZBEHANI, M., DAHLEH, M., AND MITTER, S. 2010. Dynamic pricing and stabilization of supply and demand in modern electric power grids. In *Proceedings of the 1st IEEE International Conference on Smart Grid Communications*. IEEE, 543–548.
- RURAL BROADBAND. 2012. Avanti Hylas 1 satellite broadband services. <http://www.ruralbroadband.co.uk/avanti-hylas-1-8mbps.html>.
- SACHSON, T. 2011. Subsidized data delivery via toll-free apps. eComm 2011.
- SAMADI, P., MOHSENIAN-RAD, A., SCHÖBER, R., WONG, V. W. S., AND JATSKEVICH, J. 2010. Optimal real-time pricing algorithm based on utility maximization for smart grid. In *Proceedings of the 1st IEEE International Conference on Smart Grid Communications*. IEEE, 415–420.
- SARAYDAR, C., MANDAYAM, N., AND GOODMAN, D. 2002. Efficient power control via pricing in wireless data networks. *IEEE Trans. Commun.* 50, 2, 291–303.
- SARKAR, M. 1997. Internet pricing: A regulatory imperative. In L. W. McKnight and J. P. Bailey, Eds., *Internet Economics*. MIT Press, Cambridge, MA.
- SCHATZ, A. AND ANTE, S. E. 2010. FCC chief backs usage-based broadband pricing. Wall Street Journal. December 2.
- SEGALL, L. 2011. Verizon ends unlimited data plan. CNN Money. July 6.
- SEMRET, N., LIAO, R. R.-F., CAMPBELL, A. T., AND LAZAR, A. 2000. Pricing, provisioning and peering: Dynamic markets for differentiated Internet services and implications for network interconnections. *IEEE J. Sel. Areas Commun.* 18, 2499–2513.
- SEN, S. 2011. Bare-knuckled wireless. Business Today. March 10.
- SEN, S., JOE-WONG, C., AND HA, S. 2012a. The economics of shared data plans. In *Proceedings of the 22nd Annual Workshop on Information Technologies and Systems (WITS'12)*.
- SEN, S., JOE-WONG, C., HA, S., BAWA, J., AND CHIANG, M. 2013. When the price is right: Enabling time-dependent pricing of mobile data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI'13)*. ACM, New York, 2477–2486.
- SEN, S., JOE-WONG, C., HA, S., AND CHIANG, M. 2012b. Incentivizing time-shifting of data: A survey of time-dependent pricing for Internet access. *IEEE Commun. Mag.* 50, 11, 91–99.
- SEPHTON, J. 2011a. The changing pricing model for mobile content. Strategy Analytics. November 24, <http://tinyurl.com/ContentPricing>.
- SEPHTON, J. 2011b. Spain—New Orange Delfin price plan. Strategy Analytics. July 13, <http://tinyurl.com/OrangeSpain>.
- SEPHTON, J. 2011c. Trialling mobile data—a new trend for mobile plans? Strategy Analytics. October 28, <http://tinyurl.com/mobiletrials>.

- SHAKKOTAI, S., SRIKANT, R., OZDAGLAR, A., AND ACEMOGLU, D. 2008. The price of simplicity. *IEEE J. Sel. Areas Commun.* 26, 7, 1269–1276.
- SHEN, H. AND BASAR, T. 2007. Optimal nonlinear pricing for a monopolistic network service provider with complete and incomplete information. *IEEE J. Sel. Areas Commun.* 25, 1216–1223.
- SHENG, M.-J., JOE-WONG, C., HA, S., WONG, F. M. F., AND SEN, S. 2013. Smart data pricing: Lessons from trial planning. In *Proceedings of the IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, Washington, DC, 375–380.
- SHENKER, S., CLARK, D., ESTRIN, D., AND HERZOG, S. 1996. Pricing in computer networks: Reshaping the research agenda. *Telecommun. Policy* 20, 3, 183–201.
- SHETTY, N., PAREKH, S., AND WALRAND, J. 2009. Economics of femtocells. In *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM'09)*. IEEE, 1–6.
- SMART DATA PRICING FORUM. 2012. <http://scenic.princeton.edu/SDP2012>.
- SONGHURST, D. 1999. *Charging Communication Networks: From Theory to Practice*. Elsevier.
- SUN, J. AND MODIANO, E. 2006. Channel allocation using pricing in satellite networks. In *Proceedings of the 40th Annual Conference on Information Sciences and Systems*. IEEE, 182–187.
- SVITAK, A. 2012. Demand for bandwidth drives commercial Satcom. *Aviation Week*. January 30.
- TAYLOR, T. 2000. Megaco errata. RFC 2886 (Historic). Obsoleted by RFC 3015.
- TEITELL, B. 2012. Cellphone overcharges putting a strain on many families. *Boston Globe*. September 22.
- TELECOMPAPER. 2010. Mobinil offers time-based mobile data plans. July 19, <http://www.telecompaper.com/news/mobinil-offers-timebased-mobile-data-plans>.
- THE ECONOMIST. 2009. The mother of invention: Network operators in the poor world are cutting costs and increasing access in innovative ways. Special report, September 24.
- U.S. ENERGY INFORMATION ADMINISTRATION. 2012. Wholesale market data. <http://www.eia.gov/electricity/wholesale/>.
- U.S. OFFICE OF HIGHWAY POLICY INFORMATION. 2011. Toll facilities in the United States. Publication no. FHWA-PL-11-032.
- VARAIYA, P. 1999. Pricing and provisioning of quality-differentiated services. In *Proceedings of the Information Theory and Networking Workshop*. IEEE, 38.
- VARDY, M. 2011. Telus Tuneage: Now powered by Rdio. *The Next Web*. August 3, <http://thenextweb.com/ca/2011/08/03/telus-tuneage-now-powered-by-rdio/>.
- VYTELINGUM, P., RAMCHURN, S. D., VOICE, T. D., ROGERS, A., AND JENNINGS, N. R. 2010. Trading agents for the smart electricity grid. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1*. International Foundation for Autonomous Agents and Multiagent Systems, 897–904.
- WAGNER, S., VAN DEN BERG, E., GIACOPELLI, J., GHETIE, A., BURNS, J., TAUIL, M., SEN, S., WANG, M., CHIANG, M., LAN, T., LADDAGA, R., ROBERTSON, P., AND MANGHWANI, P. 2012. Autonomous, collaborative control for resilient cyber defense (ACCORD). In *Proceedings of the IEEE 6th International Conference on Self-Adaptive and Self-Organizing Systems Workshop (SASOW'12)*. IEEE
- WAHLMUELLER, S., ZWICKL, P., AND REICHL, P. 2012. Pricing and regulating quality of experience. In *Proceedings of the 8th EURO-NF Conference on Next Generation Internet*. IEEE, Washington, DC, 57–64.
- WALRAND, J. 2008. Economic models of communication networks. In Z. Liu and C. H. Xia, Eds., *Performance Modeling and Engineering*. Springer, New York, Chapter 3, 57–90.
- WELLS, J. AND HAAS, D. 2004. *Electricity Markets: Consumers Could Benefit from Demand Programs, but Challenges Remain*. DIANE Publishing, Darby, PA.
- W-PIN. 2012. W-PIN: The first workshop on pricing and incentives in networks. <http://www.eurecom.fr/loiseau/W-PIN2012/>.
- WOLAK, F. 2006. Residential customer response to real-time pricing: The Anaheim critical-peak pricing experiment. Tech. rep., Stanford University. <http://www.stanford.edu/wolak>.
- YÁICHE, H., MAZUMDAR, R. R., AND ROSENBERG, C. 2000. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Trans. Netw.* 8, 667–678.
- YOO, C. S. 2009. Network neutrality, consumers, and innovation. *U. Chi. Legal F.* 25, 179. University of Pennsylvania Law School, Public Law research paper no. 08-40.

Received February 2012; revised October 2012; accepted February 2013