Smart Data Pricing: Using Economics To Manage Network Congestion

How can economic theory and field trials of new pricing mechanisms help find a win-win for both network operators and users?

Soumya Sen[1], Carlee Joe-Wong[2], Sangtae Ha[3], Mung Chiang[2]
[1]University of Minnesota, [2]Princeton University, [3]University of Colorado-Boulder

The past five years have witnessed a rapid evolution in pricing practices among Internet Service Providers (ISPs) in the US and other international markets, particularly in moving away from flat-rate to usage-based pricing in cellular networks [1]. In 2010, AT&T eliminated unlimited data plans and introduced a tiered plan of about $10/GB, along with throttling of heavy users and hefty overage fees. Verizon and other US ISPs have since introduced similar data plans. These changes have fueled continued debates about net neutrality and the openness of the Internet.

ISPs argue that these price increases and usage penalties are needed in response to rapid growth in data traffic, driven by increasing demand for smart devices, bandwidth-hungry applications, cloud-based services, M2M traffic, and media-rich web content [2, 3]. This explosive growth, which Cisco's VNI projects will cause an 11-fold increase in global wireless traffic between 2013 and 2018 [4], requires investment in expanding wired and wireless network capacities (e.g., by acquiring additional spectrum, deploying WiFi hotspots for offloading data traffic, investing in the backhauling infrastructure, and adopting newer technologies such as 4G/LTE and femtocells). The benefits of this capacity expansion are partly accrued by the content providers, who attract more advertising and e-commerce revenues from greater user demand and drive the demand for bandwidth even further up. ISPs therefore contend that they are trapped in a vicious cycle that does not allow them to match their prices to the costs. Measures such as throttling, data caps, and usage-based (metered) pricing are thus viewed as essential tools for regulating demand and managing network congestion.

However, simply penalizing demand would be harmful for the Internet ecosystem and might restrict network access for some users. We argue that in order to address the core issue of improving network resource management, demand growth should be seen not as a problem that will disappear if users are properly penalized, but as an opportunity for ISPs to monetize their network assets and better manage network congestion by creating the right economic incentives for their users. In other words, appreciating the role of economics in network management and understanding how these economic models can be realized in practical network systems is crucial to the Internet's long-term growth and sustainability.

Exploring this link between economic principles and network engineering has been the goal of several recent research efforts in *Smart Data Pricing* (SDP) [5]. SDP mechanisms go beyond simple byte-counting schemes to include time/location/app-based dynamic pricing, usage-based pricing, possibly with differentiated speed tiers, auction-based smart markets, WiFi offloading, proactive caching, zero-rating or sponsored content, quota-aware content adaptation, etc. In general, it asks three kinds of questions along the following dimensions:

1. *Who* should pay for bandwidth (e.g., zero-rating, sponsored content, two-sided pricing)?
2. *What* service should be charged for (e.g., transaction based pricing, quality based pricing)?
3. *How* to charge (e.g., time/location/congestion-dependent pricing, traffic offloading)?

> ***"To address the core issue of improving network resource management, demand growth should be seen not as a problem that will disappear if users are properly penalized, but as an opportunity for ISPs to monetize their network assets and better manage network congestion by creating the right economic incentives for their users."***

In this article, we focus on the question of *"How?"* by reporting on two research directions: time-dependent pricing and traffic offloading mechanisms. Although researchers have been exploring the interplay between networks and economics for many years [2, 6, 7] (for a detailed survey of various pricing proposals see [8]), the need for designing and demonstrating fully functional prototypes has become recently more urgent as user demand for data grows. Consequently, a key aspect of recent SDP research has been its focus on bridging the often-noted gap between analytical models and practical considerations:

(a) *From practice to modeling:* Network technology needs to be complemented with sound economics, but the analytical models developed should account for the real-world constraints imposed by existing technical and practical operational considerations (e.g., parameter measurability, data granularity, solution scalability and complexity, integration and deployment feasibility, regulatory requirements).

(b) *From modeling to practice:* Analytical models need to provide guidance on the economics of bandwidth pricing and policy, as well as on the architecture that will be used to implement these models in operational networks. The prototypes developed then need to be tested by deploying them in the wild.

To successfully incorporate this interplay in SDP, researchers need to take a holistic approach that brings together ideas from economics, networking, information systems, human-computer interaction (HCI), etc. In general, SDP research has three stages, each with interesting research questions that are essential components of realizing a new pricing scheme. For example:

(1) Analytical modeling – How can we create economic models for computing optimized prices or incentives that ISPs will be willing to offer and users will be ready to accept for modifying their bandwidth consumption behavior in the desired manner?

(2) System development – How can we develop scalable systems and related signaling protocols for offering these incentive mechanisms? How should the various required functionalities (e.g., congestion measurement, traffic profiling, delay-optimal scheduling) be divided between the network backend and the end-user devices?

(3) Field trials – How can we use principles from HCI to design end-user interfaces so that users can understand and respond to the pricing signals?

A series of new initiatives, such as the IAB's Workshop on Internet Technology Adoption & Transition (ITAT) [8], Smart Data Pricing Forum [5], IRTF WG on the Global Access to the Internet for All (GAIA) [10], etc., are providing momentum for such interdisciplinary collaborations. This article discusses two complementary research themes in SDP – time-dependent pricing (TDP) and traffic offloading – that aim to reduce network congestion by providing users with incentives and mechanisms to "shift" their usage to less congested times or to other supplementary networks. We also consider SDP's implications for the Internet's long-term sustainability and accessibility to a wider user population.

> ***"To successfully incorporate models and practice in Smart Data Pricing, researchers need to take a holistic approach that brings together ideas from economics, networking, information systems, human-computer interaction, etc."***

Case 1: Time-Dependent Pricing

Much of the need for expanded network capacity is due to large peak demands created by users' simultaneous consumption of data; Cisco's VNI predicts that peak hour traffic will grow at 64% CAGR [4]. Yet as Odlyzko *et. al.* [11] correctly point out, ISPs' attempts to slow this growth by transitioning from flat-rate to usage-based pricing are unlikely to solve this problem. To discourage large peak demand, prices need to have a temporal component, *i.e.,* they should vary over different times of the day, as in time-dependent pricing. Only then will users be incentivized to spread out their demand over time, thus improving network resource utilization by reducing the peaks and filling up the valley periods. Writing for Google's Public Policy blog, Vinton Cerf advocated a similar view - "*Network Management also should be narrowly tailored, with bandwidth constraints aimed essentially at times of actual congestion.*" Cerf cautions against ISPs' rush to change their pricing, favoring a more detailed study on the efficacy of such dynamic or time-varying pricing mechanisms. The work we present here is one such attempt at understanding the efficacy and feasibility of TDP for mobile data.

The telecommunications industry has long practiced dynamic TDP plans for voice calls to respond to demand variability in call volume by adjusting users' prices/incentives. However, dynamic pricing plans for data traffic, in spite of their theoretical potential to make resource allocation much more efficient, have remained largely unrealized in the global market, which may be partly due to the gap between the large body of analytical works on this topic and the lack of functional prototypes implementing these ideas.

The academic literature on dynamic pricing theories is quite extensive [3] and includes responsive pricing (which sets prices so as to keep user demand under a certain threshold), proportional fairness pricing (which sets prices to optimize a proportional fairness criterion on the amount of bandwidth allocated to different users), priority pricing (which explicitly accounts for QoS by allowing users to pay less by accepting a longer delay at congested times), and "smart market" auction pricing (which decides whether to admit a packet into the network at congested times based on the user-specified bid attached to that packet). But such congestion pricing schemes face two practical challenges: (a) users prefer flat-rates over the uncertainties associated with near real-time fluctuations in prices [7, 12], and (b) users are often found reluctant to delegate price bidding or traffic scheduling to automated agents, desiring the psychological assurance of manual control despite automation's greater convenience [12]. For users to be comfortable with dynamic pricing for data, ISPs have to provide some guarantees on available future incentives or prices, design intuitive user interfaces to aid manual decision-making, and demonstrate the underlying system's feasibility with a proof-of-concept prototype.

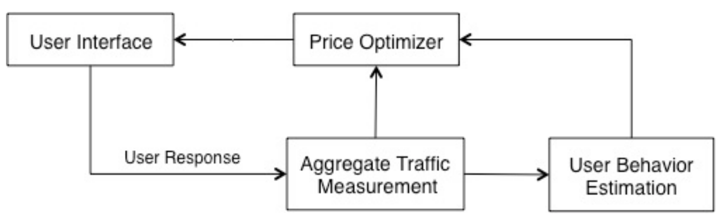

Figure 1a: Control-feedback loop of dynamic time-dependent pricing.

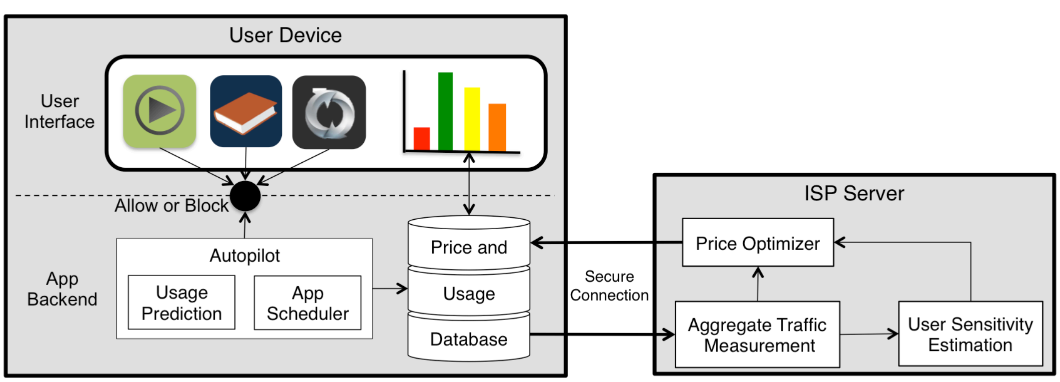

Figure 1b: Functionality separation between the user-side and operator-side devices.

Practice to Models: To address the previous practical issues within an analytical framework, we have explored an alternative pricing approach – dynamic day-ahead time-dependent pricing [13].

In this scheme, the ISP computes hourly prices one day in advance and advertises them to all users. The provider continues to compute new prices to maintain a sliding one-day window of announced prices, *i.e.*, a new price point for the 24^{th} hour is computed and announced every hour. Users thus receive day-ahead price guarantees while being rewarded for shifting some portion of their data usage (e.g., non-critical traffic) according to these announced price points. The ISP can measure changes in usage volume at different hours of the day in response to the given set of prices and estimate users' willingness to shift different types of traffic, which in turn is used to compute the next set of optimized prices. These prices are computed using a convex optimization formulation that minimizes the provider's total cost of overshooting capacity and the cost of providing these incentives to users, while accounting for current estimates of users' willingness to shift traffic, temporal variations in usage volume, and capacity constraints [13]. Thus, the framework requires a control-feedback loop between the ISP and users (refer to Figure 1a).

In addition to considering users' psychological preference for certainty regarding the future price points, TDP must accommodate technological and regulatory concerns. Firstly, in order to preserve users' privacy, our formulation does not require the ISP to track each individual user's usage pattern, precluding the need for any Deep Packet Inspection (DPI) to realize this pricing scheme. Secondly, the formulation remains computationally efficient as the number of network users grows: the model in [13] implicitly assigns the aggregate traffic from all users and applications into virtual traffic classes that are characterized by different delay and price sensitivity estimates. It then accounts for heterogeneity across users by separately modeling the probabilistic deferral behavior for each traffic class. Thirdly, this optimization model avoids the use of utility functions, which can be hard to measure quantitatively. Instead, all parameters can be either directly measured (e.g., changes in usage volume in each period in response to prices) or estimated (e.g., users' price and delay sensitivities for different traffic classes) by the ISP without compromising the system's scalability or users' privacy.

Models to Practice: The system used to realize this day-ahead TDP for mobile data is shown in Figure 1b. The aggregate traffic measurement, user sensitivity estimation, and price computation engines reside on the ISP side, while a mobile application that communicates with the pricing engine resides on each end-user device. The primary purpose of the mobile application was to provide users with information on available price points, but it also had optional features like usage monitoring and alerts, as well as an auto-pilot mode for automated scheduling of applications based on the available prices and user-specified delay sensitivities. For privacy reasons, such scheduling information remains in the app and is not communicated to the ISP. Additional details regarding the UI design, efficacy, and user response to these features are available in [11].

A prototype of this system was developed and tested for eight months in different phases of a randomized field experiment in Princeton, New Jersey, with 50 mobile users of a very large US wireless service provider. During the trial, we effectively became a resale ISP offering this data plan with day-ahead dynamic prices to the trial participants. We separated the 3G traffic of these participants from that of other customers using an Access Point Name (APN) setup, which tunneled participants' 3G traffic from the ISP's core network into lab servers, as shown in Figure 2. The participants installed the TDP mobile application on their iOS devices. WiFi usage, voice calls, and SMS were not included in the trial traffic as these do not count towards 3G data caps.

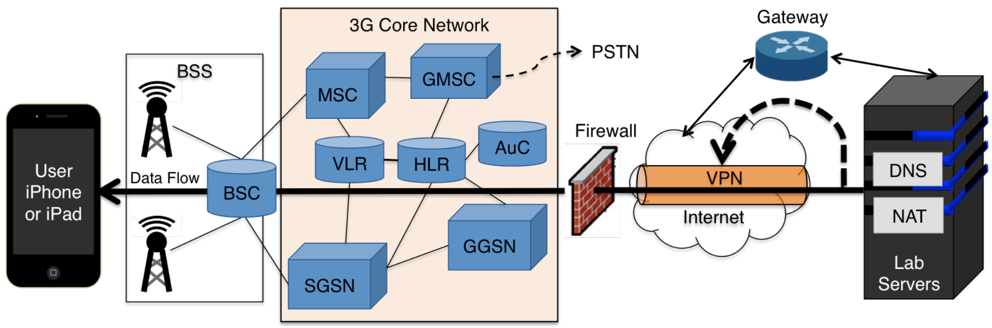

Figure 2: Field trial setup for dynamic day-ahead time-dependent usage-based pricing.

The mobile application running on users' devices displayed the prices/discounts available for the next 24 hours in a color-coded format as shown in Figure 3(a). Each price is color-coded by its discount rate, e.g., red (0-10%), orange (11-19%), yellow (20-29%), green (>30%). Users could view their usage history superimposed over the offered prices to visualize how much they spent and saved on data. Additionally, to help users save money, the app provided interfaces for users to see their top 5 bandwidth-consuming applications, set alerts, and configure their weekly budgets and app delay sensitivities for automatic scheduling, as shown in Figures 3(b)-(e).

Figure 3 (a): Landscape view of superimposed price and usage history by day, week, and month,

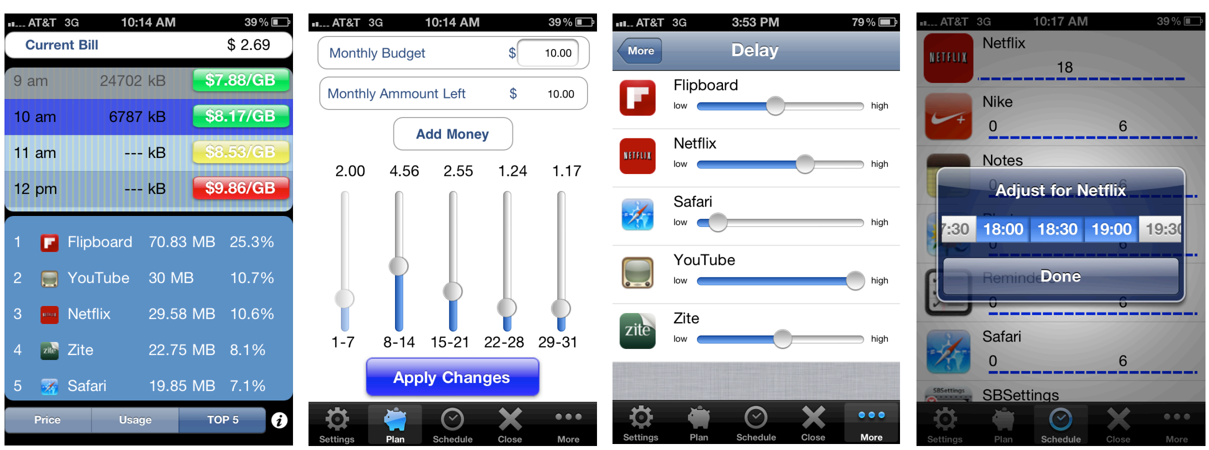

(b) (c) (d) (e)

Figure 3: (b) View of the top-5 bandwidth consuming apps in the bottom split-screen, (c) Weekly budget adjustment screen, (d) App-delay sensitivity settings screen, and (e) App-specific temporal blocking in parental control.

Using the analytical model in [13], we offered optimized day-ahead time-varying price discounts on the baseline price of $10/GB to all users. We found that users did shift their traffic from high-price to low-price periods under TDP: the overall usage of the iPad users decreased by by 10.1% in high-price periods and increased by 15.7% in low-price periods, that is, for most users, the average usage decreased in high-price periods relative to average usage at the same hours before the trial [13]. Additionally, by focusing on usage in consecutive periods where the discounts differed by only 1% but the colors of the price indicator bars were different (*e.g.,* comparing usage volumes in a yellow price period with a 29% discount to the following green price period with a 30% discount), we found a significant change in usage even though the absolute % change in discount was only 1%. This indicates that users paid more attention to the color-coding rather than the actual value of the price discounts. This emphasizes the need for careful and intuitive user interface design to ensure that users are able to understand and respond to pricing signals [12].

With the optimized day-ahead time-dependent prices, resource utilization at off-peak hours nearly doubled, indicating that TDP can also be effective in improving the utilization of network capacity by flattening and distributing the demand over different times of the day. The resulting maximum observed daily peak-to-average ratio (PAR) decreased by 30% with dynamic day-ahead TDP, and approximately 20% of the PARs from the pre-trial period were larger than the maximum PAR with TDP. Although the population size of this field trial is somewhat limited due to the complexity of conducting such trials in academic settings, the results are promising: they demonstrate that it is possible not only to operationalize dynamic pricing for mobile data but also to use time-varying prices to change users' behavior.

It also bears mention that TDP can also be reflected as time-dependent sponsored content, whereby consumers do not see any price fluctuations, but valley discount is reflected in the price that sponsoring parties like content providers or enterprises pay to an ISP.

Case 2: Traffic Offloading

In addition to shifting demand from peak to valley time periods, ISPs can alleviate congestion by shifting demand off of their cellular networks onto supplementary networks like WiFi or femtocells. This process is also known as traffic offloading. Many ISPs have started to encourage offloading by selling bundles of base and supplementary technologies (e.g., Orange offers a £2 bundle of 3G and WiFi hotspot access) [14]. Conventional wisdom says that increasing the access price of the base technology (*i.e.,* 3G, 4G) will encourage users to purchase the bundle and offload more to the supplementary network. This seems to be a direction that many service providers have been pursuing with various penalty mechanisms on their base technology. But this strategy does not account for the fact that the supplementary network can itself become more congested as more and more users offload their traffic, potentially making it less attractive to users. In other words, there is a complex interaction between the prices of the base and supplementary technologies, the relative network congestion externalities of these two technologies, and the coverage area of the supplementary technology. Economic models and their practical realization in field trials can help ISPs design more effective offloading mechanisms. In this section, we discuss some related results from our studies that focus on the theory as well as implementation of ideas that can improve offloading performance and make it easier for users to make such offloading decisions.

***Practice to Models*:** Understanding how users will decide to adopt the base technology's network or a bundle of base and supplementary technologies, as well as deriving the resulting equilibrium and transient market outcomes, requires analytical models that incorporate practical issues like congestion on both networks and the coverage area of supplementary networks. Sen *et. al.* [15] introduced a model to study the dynamics of competition between two generic network technologies with cross-network externalities in the presence of network gateways or converters.

In [14], we extend this framework to develop an analytical model in which users individually make their adoption decisions based on several factors, such as the technologies' intrinsic qualities, users' heterogeneity in the evaluation of these qualities, negative congestion externalities from the presence of other subscribers on the technologies, and the access rates charged by an ISP.

Using the analytical model introduced in [14], we study how user-level decisions translate into aggregate adoption dynamics and characterize the equilibrium outcomes for different system parameters. The model reveals that seemingly intuitive strategies can sometimes have unintended consequences for congestion on the base network technology: for example, increasing the coverage area of the supplemental technology can increase traffic on the supplemental network, inducing some users to drop the bundled service altogether and use only the base technology. Similarly, increasing the base technology's access price can cause some users to find the bundle to be too expensive (i.e., its coverage and relative network congestion conditions do not offset the decrease in some users' utility from the base price increase), and so they will drop the bundled service and just use the base technology, contributing to an increase in congestion on that network. Careful analysis using economic models that capture all these different aspects can therefore prove very useful in providing guidance and insights into potential outcomes from network pricing and policy changes.

Models to Practice: While analytical frameworks like the one discussed above can provide insight into users' long-term adoption behaviors (*i.e.,* equilibrium, stability etc.), users' minute-to-minute offloading decisions reflect their more immediate concerns. Users face a three-way tradeoff between cost, throughput, and delay: while they can save money and receive a higher throughput by waiting for WiFi access, they may not want to wait for some critical applications. To navigate this tradeoff, we propose a practical, cost-aware WiFi offloading system called AMUSE (Adaptive bandwidth Management through USer-Empowerment) [16]. AMUSE is a user-centric tool that learns the user's behavior and mobility pattern to decide which applications to offload to what times of the day, thus enabling users to stay within their data caps. In doing so, it uses a utility optimization algorithm to decide if and how much 3G bandwidth needs to be allocated to each application at any given time based on the user's available budget, application delay sensitivities, and inputs from prediction algorithms regarding the user's demand patterns and WiFi availability in the future.

To implement this allocation in practice, AMUSE uses a receiver-side TCP bandwidth control algorithm that enforces each application's assigned download rate by controlling the TCP advertisement window on the user side. Thus, this algorithm is fully contained on end-user devices and does not require any modification on the TCP server side, making it suitable for real-world deployment. Figure 4 shows these different modules of the AMUSE system and their interaction.

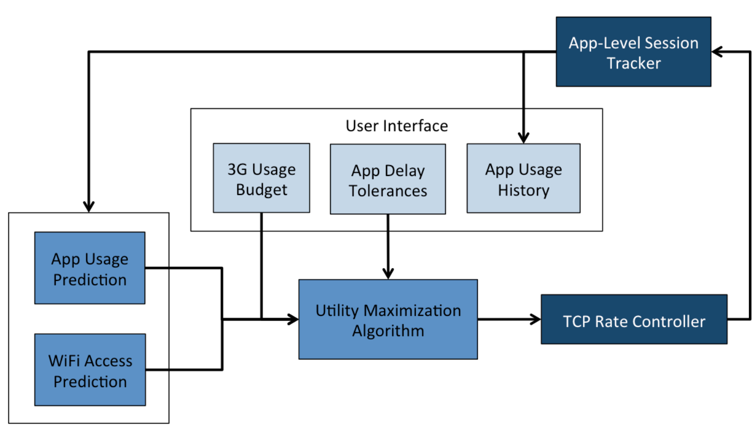

Figure 4: System modules and interaction across various components of AMUSE.

An AMUSE system implementation on Windows 7 tablets, when tested with simulated user behavior based on 3G and WiFi usage and availability data collected from a field study of 37 mobile users, showed that other offloading algorithms yield 14% and 27% lower user utilities than AMUSE for light and heavy users, respectively [16]. Intelligently managing users' competing interests for cost, throughput, and delay via user-side management tools can therefore facilitate user decisions, leading to offloading benefits that could be even further improved when integrated with ISP-centric approaches (e.g., ad-hoc authorization for offloading to hidden WiFi access points, or market based solutions for offloading from base stations to third-party owned WiFi access points and femtocells [17-19]). Such measures to push some control from the network core out to the end users complement previous and existing efforts of the networking research community [20].

"These two approaches to managing network congestion, time-dependent pricing and delay optimal traffic offloading, can help create a "win-win" solution for both the ISPs and their users."

Creating a "Win-Win" Solution to Bridge the Digital Divide

The two approaches to managing network congestion presented earlier, namely realizations of time-dependent pricing and delay optimal traffic offloading, can help create a "win-win" solution for both ISPs and their users. ISPs benefit by reducing peak congestion, while users are offered more choices and technologies to save on their monthly bills.

Research on network resource pricing also has implications for bridging the digital divide between those who can and who cannot access the Internet regularly. Rural local exchange carriers (RLECs) often suffer from congestion in their wired networks due to the persistence of the middle-mile problem – although the cost of middle mile bandwidth has declined over the years due to an increase in the demand needed to fill the middle mile, the bandwidth requirements of home users have also increased sharply. The cost of middle mile upgrades to meet the FCC's target speed of 4 Mbps will be substantial and is a barrier to digital expansion in rural areas [21]. New access pricing mechanisms like TDP can help bring down middle mile investment costs by reducing RLECs' peak capacity provisioning or leasing needs and improving resource utilization in the valley periods. Thus, providers can match their prices to the cost of delivery while also creating incentives for light users to adopt broadband services. Instead of being charged by the volume of data consumed, users can have the option of saving on their monthly bill by choosing "when" to consume this data. An extension of this idea can be used to create ultra-affordable data plans for delay-tolerant users that allow automated opportunistic access to the network only when large discounts are available, that is, when the network is lightly loaded. Implementation of such dynamic time-dependent pricing and opportunistic offloading schemes can help service providers to better utilize their available resources and increase Internet adoption by being financially attractive to light data users.

Content providers have also made efforts to bridge the digital divide with zero-rating, app-based pricing, and sponsored content [22], all of which open up further interesting questions for SDP research, for example, should open platform for sponsored content be adopted to be compatible with net-neutrality? How will such subsidized plans impact broadband adoption and network congestion? These are the questions that much of the current SDP research is focused on.

"Pricing of network services has implications not only for managing available network capacity but also for bridging the digital divide."

Conclusions

As the demand for bandwidth grows in wireless and wired networks, Internet service providers are pursuing penalty mechanisms like throttling, capping, overage charges, usage-based fees, etc., to manage their available network capacity. But such measures are arguably suboptimal and even harmful to the Internet ecosystem. In contrast, ideas from economics can help design incentives and pricing policies that will be beneficial to both service providers and their users. Although many analytical models for pricing-based network management have been proposed, their implementation in practice has only started recently. Tackling today's growing challenges requires not only developing analytical models that incorporate practical concerns (*e.g.,* measurability, scalability, privacy, user behavior) but also demonstrating their efficacy and feasibility through prototypes and field trials.

In this article, we focused on the mechanism of "*shifting*" demand with two complementary efforts that aim to alleviate network congestion by creating incentives and mechanisms to modify user behavior, *i.e.,* shift demand either to less congested times (using time-dependent pricing) or to a supplementary network (using delay-optimized traffic offloading). The results indicate that such measures, if implemented with intuitive designs, can help ISPs better monetize and manage their network capacity, while empowering users with more options to avoid hefty overage fees. With implicit pricing signals and automated decision-making, these solutions can be readily adapted to emerging M2M and IoT applications. These approaches also have positive implications for making Internet access affordable for users, thereby contributing to efforts in bridging the digital divide.

SDP, as mentioned in the introduction, has a growing scope and is much broader than the two particular cases that we have space to outline in this article. The complementary questions of "*Who*" and "*What*" are gaining traction in addition to the question of "*How*" on which we focused here. Going forward, researchers will need to move beyond models, trials, and testing of prototypes and start considering integration with existing infrastructure, the design of pricing and signaling protocols, and regulatory concerns. Smart Data Pricing research will help bring together academics, network providers, content providers and the e-commerce industry, and policy makers to design and deploy new mechanisms that will help ensure the sustainable growth of the Internet, mobile, and content markets.

Acknowledgements

We would like to thank many collaborators on SDP in academia and industry, especially to Dr. Krishan Sabnani of Bell Labs for his encouragement and feedback on this article.

Bibliography

[1] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. *Incentivizing time-shifting of data: A survey of time-dependent pricing for Internet access.* IEEE Communications Magazine, vol. 50, no. 11 (2012), 91–99.

[2] El-Sayed, M., Mukhopadhyay, A., Urrutia-Valdés, C., Zhao, Z. J., *Mobile Data Explosion: Monetizing the Opportunity through Dynamic Policies and QoS Pipes*, Bell Labs Technical Journal, vol. 16(2), pp. 79-100, 2011.

[3] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. Smart Data Pricing: Economic solutions to network congestion. *ACM SIGCOMM eBook on Recent Advances in Networking,* 2013.

[4] Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013—2018, Cisco, 2014, Available at https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.pdf

[5] Smart Data Pricing Forum. www.smartdatapricing.org

[6] Gupta, A., Stahl, D. O., and Whinston, A. B. *The economics of network management.* Communications of the ACM, 42(9), (1999), 57-63.

[7] Odlyzko, A. Will smart pricing finally take off? *Smart Data Pricing*, eds. Sen. S, Joe-Wong, C., Ha, S., and Chiang, M., Wiley, 2014.

[8] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. *A survey of broadband data pricing: Past proposals, current plans, and future trends*, ACM Computing Surveys, vol. 46(2), (2013), 1-37.
[9] Lear E., Report from the IAB workshop on Internet Technology Adoption and Transition (ITAT), IETF Network Working Group, Internet-Draft, Available at http://tools.ietf.org/html/draft-iab-itat-report-03
[10] Sathiaseelan, A, Global Access to the Internet for All (GAIA), IRTF Working Group, Information Society, 2014.
[11] Odlyzko, A., St. Arnaud, B., Stallman, E., and Weinberg, M. *Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service*, Public Knowledge, May 2012.
[12] Sen, S., Joe-Wong, C., Ha, S., and Chiang, M. *When the price is right: Enabling time-dependent pricing of broadband data,* Proc. of ACM SIGCHI (2013), 2477-2486.
[13] Ha, S., Sen, S., Joe-Wong, C., Im, Y., and Chiang, M. *TUBE: Time-dependent pricing for mobile data.* ACM SIGCOMM Computer Communication Review, 42(4) (2012), 247-258.
[14] Joe-Wong, C., Sen, S., and Ha, S. *Offering Supplementary Network Technologies: Adoption Behavior and Offloading Benefits.* IEEE/ACM Trans. on Networking (doi: 10.1109/TNET.2014.2300863), 2014.
[15] Sen, S., Jin, Y., Guerin, R. and Hosanagar, K. *Modeling the dynamics of network technology adoption and the role of converters.* IEEE/ACM Transactions on Networking, 18(6) (2010), 1793-1805.
[16] Im, Y., Joe-Wong, C., Ha, S., Sen, S., Kwon, T., and Chiang, M. *Mobile offloading decisions: Cost, throughput, and delay tradeoffs.* Proc. of IEEE INFOCOM, pp. 435-439, 2013.
[17] Iosifides, G., Gao, L., Huang, J., and Tassiulas, L. *A Double Auction Mechanism for Mobile Data Offloading Markets,* IEEE Transaction on Networking, (accepted) July 2014.
[18] L. Gao, G. Iosifidis, J. Huang, L. Tassiulas, and D. Li, *Bargaining-based Mobile Data Offloading*, IEEE Journal on Selected Areas in Communications, vol. 32, no. 6, pp. 1114 - 1125, June 2014.
[19] M. Cheung and J. Huang, *Optimal Delayed Wi-Fi Offloading*, IEEE WiOpt, Tsukuba Science City, Japan, May 2013.
[20] Briscoe, B., Darlagiannis, V., Heckman, O., Oliver, H., Siris, V., Songhurst, D., and Stiller, B. *A market managed multi-service Internet*, 2003. http://www.m3iproject.org/.
[21] Glass, V., Stefanova, S., Dibelka, R. Customer Price Sensitivity to Broadband Service Speed: What are the Implications for Public Policy? *Smart Data Pricing*, eds. Sen. S, Joe-Wong, C., Ha, S., and Chiang, M., Wiley, 2014.
[22] Andrews, M., Rieman, M., Wang, Q., and Ozen, U. *Economic Models of Sponsored Content in Wireless Networks with Uncertain Demand,* 2nd Workshop on Smart Data Pricing, 2013.