SpeedGate: A Smart Data Pricing Testbed Based on Speed Tiers

Yih-Farn Chen*, Rittwik Jana*
*AT&T Labs-Research, 180, Park Avenue, Florham Park, NJ, 07932, USA

Abstract — The explosive growth of cellular traffic and its highly dynamic nature often make it increasingly expensive or even infeasible for a cellular service provider to provision enough cellular resources to support the peak traffic demands. Some service providers have started exploring various economic incentives, including smart data pricing, to manage network congestion. We present SpeedGate, a smart mobile data pricing testbed that allows a service provider to experiment with different dynamic pricing strategies. SpeedGate maintains persistent VPN connections to smartphones as users roam between different wireless networks (3G, 4G/LTE, WiFi). The maximum available bandwidth per user session can be adjusted according to various data pricing strategies. We report preliminary results on two trials with a total of 29 users for assessing their willingness to pay (WTP) for various speed tiers. Preliminary observations suggest the challenges of QoS guarantees through speed tiers in the field, the limited dynamic range of WTP values from individual users for different speed tiers, and potential opportunities for auction-based dynamic pricing.

I. INTRODUCTION

As mobile data traffic has been growing explosively over the past few years and is expected to grow 18-fold between 2011 and 2016 [3], how to provide high-quality mobile Internet services to satisfy the ever-increasing traffic demands is becoming an urgent issue faced by today's wireless service providers. The nature of this problem is different from traditional resource allocation problems in which resources are under the control of a single service provider. The smartphone users have the choice to access a large array of third-party services with varying bandwidth demands at any time and almost from anywhere. How to incentivize data consumers to adjust their usage behaviors to help manage the network congestion is a significant challenge that service providers need to address. Congestion caused by the onslaught of popular, data-intensive smartphone apps can lead to a poor user experience, including slow access to content and dropped call. One answer has been to add more cellular resources, at considerable cost, and to complement cellular capacity with other technologies, such as WiFi and Femtocells. While these solutions have successfully reduced network congestion and offloaded traffic from the cellular network, there is still work to be done. As one step toward keeping up with the exponential data growth and solving the congestion problem, both researchers and service providers are exploring economic incentives such as smart data pricing.

Unfortunately, it is difficult for a service provider to predict the impact of pricing on the network congestion level, revenues, profits, and/or social welfare without conducting a large-scale trial, which may require significant engineering efforts in the wireless packet core, especially in the policy/QoS components. Figure 1 shows a simplified view of the logical policy charging and enforcement architecture of a cellular system.[1] The 3GPP PCRF (Policy Charging and Rules Function)[12] is a new logical element in the 3GPP packet core that makes policy/QoS decisions based on various inputs and pushes the policies that are to be enforced for a subscriber's data session down to the PCEF (Policy and Charging Enforcement Function) component. PCRF also serves as the interface with external systems for any required subscriber information to allow for the necessary policy decisions to be made. The GGSN provides external connectivity between user equipment (UE) and external packet data networks, while PCEF sits on the data path to provide gating control, i.e., the blocking or allowing of packets based on the QoS policies. While the PCRF/PCEF architecture provides some of the key building blocks for smart data pricing and allows volume-based charging, time-based charging, event-based charging, etc, it requires significant engineering efforts to modify the wireless core to support various dynamic pricing trials. Our testbed aims to provide a capability of running pricing trials with a large number of participants without accessing or modifying the PCRF/PCEF components in the core network.

There has been a lot of research in pricing models for the Internet. They can be grouped in mainly two categories, namely static and dynamic. Static plans charge users according to predetermined rates. They are easier to understand for the end user; however, these plans do not reflect users' willingness to pay (WTP) or network conditions in any way [5], [6], [10]. Dynamic pricing, on the other hand, tries to change the user prices according to the network conditions [9], [14], [8], [15]. We recognize the importance of such dynamic pricing plans in the context of mobility. There is little evidence of a pricing testbed that has been developed to try out different pricing strategies. Sen et al. [8] recently proposed such an architecture (TUBE), which is mainly designed for time-delayed pricing (TDP).

In this paper, we describe an alternative smart data pricing (SDP) testbed called SpeedGate that has several unique features: (a) **carrier-independence:** It allows mobile users from

[1]UTRAN stands for UMTS Terrestrial Radio Access Network, SGSN stands for Serving GPRS Support Node, GGSN stands for Gateway GPRS Support Node, and HLR/HSS stands for Home Location Register/Home Subscriber Server.

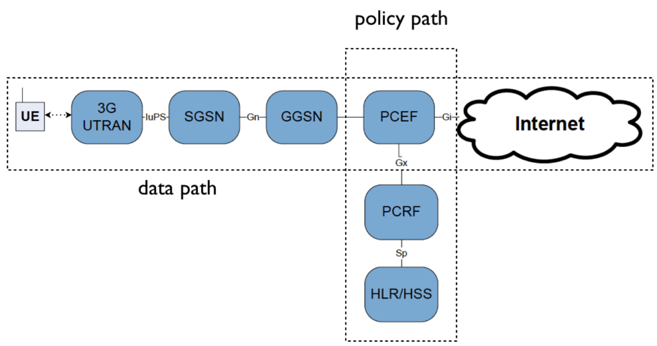

Fig. 1. A Simplified Architecture Overview of the 3GPP PCRF and PCEF functions

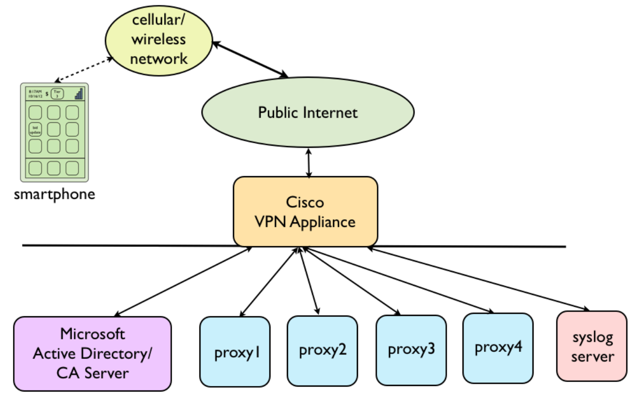

Fig. 2. SpeedGate: A Smart Data Pricing Testbed with Speed Tier Proxies

any carrier and anywhere to bring their own smartphones or tablets to participate in the pricing trials. (b) **persistent connection:** The testbed maintains a persistent VPN connection even when the subject hops among different wireless networks (3G, 4G/LTE, WiFi). (c) **dynamic speed tier assignments**: The testbed can redirect smartphone users to different ports on proxy servers in real time and adjust their maximum QoS level (allowed bit-rates) dynamically based on the smart data pricing algorithms. We have used SpeedGate to quantify users' utilities, what really affects users' decisions and how malleable users are towards different pricing strategies, all without significant engineering efforts in the carriers' wireless packet core. The rest of the paper is organized as follows: Section II describes the hardware and software components of our SDP testbed SpeedGate. Section III describes two trials that we ran on our SDP testbed to measure users' willingness to pay (WTP) for a given quality of service. The first trial uses the testbed to provide different speed tiers to cellular data users. The second trial uses the same testbed (specifically bandwidth proxies) in a WiFi environment. The second trial also provided a survey-based economic technique for estimating the valuation of user experience. This type of technique has been used in experimental economics, known as contingent valuation [7]. Section IV describes how SpeedGate can be extended with an auction engine to support dynamic pricing based on Generalized Second Price (GSP) Auction. Section V concludes with a summary and future work.

II. Testbed

Our smart data pricing (SDP) testbed is designed to satisfy one key requirement in conducting dynamic pricing trials: the testbed must be able to dynamically adjust the QoS level (expressed as the maximum bit rate) of each individual trial user. This section describes the hardware and software modules that have been developed to meet this requirement and to facilitate running trials using various smart data pricing algorithms. To date we have used this infrastructure to estimate customers' willingness to pay (WTP) for given speed tiers. In the future, we will leverage the same infrastructure for conducting a trial based on Generalized Second-Price Auction (GSP)[2]. We highlight the salient features that are necessary to conduct such a trial.

Figure 2 shows the architecture of the SpeedGate testbed, which can be configured to simulate multiple speed tiers (we have used speeds ranging from 128kbps to 32Mbps during our various trials). It allows us to dynamically switch any smartphone user to a different speed tier based on the dynamic pricing algorithms. A smartphone owner needs to first download the Cisco AnyConnect VPN client[1], obtain credentials from the trial operators, and then establish the VPN connection to the testbed. The Cisco VPN appliance[4] can maintain a VPN connection persistently even when a smartphone switches among LTE, 4G, 3G, or various WiFi networks. The Microsoft Active Directory/Certificate Authority server issues a certificate for each smartphone during the initial connection. The certificate instructs the VPN appliance which proxy server and port number to redirect the traffic of all future VPN connections from this smartphone to. It also allows each smartphone to authenticate itself to the VPN appliance as it reestablishes the VPN connection. Dummynet traffic shaper (through ipfirewall)[13] is installed on each proxy server to throttle the TCP connection at different speeds for different ports. Currently, each proxy server can support up to 64 smartphone VPN connections and therefore, the testbed can only support up to 256 concurrent users. We have plans to move the whole infrastructure to the cloud to scale the testbed to handle a larger trial. The syslog server is set up to collect logs on the connection times (but not the details of the traffic) of each VPN user.

Note that our approach is different from Paris Metro Pricing[11], where the network is partitioned into several logically separate channels with the only difference being the prices paid for using them. Channels with higher prices would attract less traffic and thereby provide better service. Paris Metro Pricing uses price as the only primary mechanism to manage traffic and there is no QoS guarantee. In addition, the scheme may sacrifice some of the utilization efficiency of the network. Our platform is intended for dynamic pricing schemes such as GSP[2]. In the GSP scheme, a user is allowed to bid at a price of his or her choice and then assigned a speed tier dynamically based on the current resource constraints and bids from other users. The throughput is close to the speed tier as long as the underlying wireless speed is above the speed tier. The testbed can dynamically adjust the number of ports available for different speed tiers to either optimize the network utilization, the revenues, the total social welfare, or to follow the PFauc (Proportional Fairness with Auction) algorithm proposed in the GSP paper[2].

The SDP testbed can be used for various dynamic pricing trials without altering the core network of a carrier's wireless infrastructure. In addition, trial participants can use any carrier (and any of a large number of smartphones supported by the VPN appliance). In the next section, we first describe two specific trials we conducted on the platform to measure the willingness to pay (WTP) of users for various speed tiers, followed by a discussion on how the testbed may be used to conduct a trial on GSP for Congestion Pricing[2] with the addition of an auction manager.

III. Assessing a user's Willingness To Pay (WTP): Trials and Evaluation Results

We now describe two trials ran on our SDP testbed for assessing a user's willingness to pay for a given quality of service. The quality of service can entail different speed tiers, delay-bandwidth product, etc.

First trial: Measuring WTP through two 24-hour treatments in the field

The trial entailed carefully recruiting students from multidisciplinary faculties. We observed the user's WTP for different allocated speed tiers. For the purpose of anonymity, we will not report on the particular speeds used, instead, we will declare the different allocated speeds as tiers 1 to 4. Higher tier is allocated a higher speed. Each user will experience two speed tiers in two separate treatments (higher to lower speed tier, or lower to higher speed tier). The users are required to bring their own smartphones and download the Cisco VPN AnyConnect client first before the treatment starts.

Using the testbed described in section II, we performed a preliminary trial with 17 students.

- After initial screening and provisioning of users into the system, users are first subjected to treatment A for 24 hours. The 24-hour treatment period was necessary to ensure that there is sufficient time for the subject to exercise the applications that he/she uses on a daily basis.
- They are then requested to report their WTP for the speed tier they experienced in that treatment after receiving an SMS message sent to their phones. The subjects are also told the current typical wireless pricing plan as a basis for comparison.
- Upon successful receipt of their responses, the users are migrated to a new speed treatment B for an additional 24 hours.
- Users are once again requested to state their WTP for treatment B after 24 hours.

Table III shows the treatments applied to the 17 users and various responses from the users. Some users did not reply in time after 24 hours and were therefore not subject to the subsequent treatment. This explains why some users' WTP values are missing. Notions like "p1, 3010, tier3" in the second and fourth columns state that the subject's VPN connection was re-directed to proxy server 1 on the SDP testbed, port 3010, which uses speed tier 3. There are a couple of interesting observations that we make.

TABLE I
WTP parameters and initial trial results

User	Treatment A	WTP A$	Treatment B	WTP B $
stu10	$p1, 3010, tier3$	0.00	$p1, 3002, tier1$	0.00
stu11	$p1, 3011, tier3$	20.00	$p1, 3003, tier1$	20.00
stu12	$p1, 3012, tier3$		$p1, 3004, tier1$	
stu13	$p1, 3013, tier4$	20.00	$p1, 3005, tier2$	25.00
stu14	$p1, 3014, tier4$		$p1, 3006, tier2$	
stu15	$p1, 3015, tier4$	40.00	$p1, 3007, tier2$	70.00
stu16	$p1, 3016, tier4$		$p1, 3008, tier2$	
stu17	$p1, 3005, tier2$	10.00	$p1, 3001, tier1$	
stu18	$p1, 3006, tier2$	10.00	$p1, 3002, tier1$	10.00
stu19	$p1, 3007, tier2$	25.00	$p1, 3003, tier1$	25.00
stu20	$p1, 3008, tier2$	10.00	$p1, 3004, tier1$	10.00
stu21	$p1, 3001, tier1$	5.00	$p1, 3005, tier2$	3.00
stu26	$p1, 3002, tier1$	0.00	$p1, 3010, tier3$	
stu27	$p1, 3003, tier1$	15.00	$p1, 3011, tier3$	15.00
stu28	$p1, 3004, tier1$	15.00	$p1, 3012, tier3$	5.00
stu29	$p1, 3005, tier2$	5.00	$p1, 3013, tier4$	15.00
stu30	$p1, 3006, tier2$	0.00	$p1, 3014, tier4$	

- A large number of users kept their WTP amount similar in both treatments: Stu11, 18, 20 and 27 did not adjust their WTP for different treatments. This could have been explained by a variety of factors. We did not impose a set of tasks to be completed (i.e. users were free to use their phones as they wish). A user who typically checks emails a couple of times a day will not be affected drastically by a speed tier change. Retrieving emails does not require a large amount of bandwidth within strict timelines.
- On the other hand, a user who consumes a lot of video data will have very different expectation. Interruption-free viewing is critical for such a user. Stu29 for example, increased his WTP significantly going from a lower speed to a higher speed.
- stu10 switched to campus WiFi completely (as he stated in his response) and opted to pay for $0 for either speed tier. This is probably true for stu26 as well.
- What surprised us the most was when a subject opted to give a lower price for a higher speed tier. This happened in the cases of stu13, stu15, stu21, and stu28. As we stated previously, while the SDP testbed strives to provide the maximum throughput specified by the speed tier, due to the varying wireless coverage the subject may experience, there is really no guarantee that they will indeed receive that speed. In general, it is risky for any wireless carrier to guarantee a particular speed for the same reason - and the user may find it annoying when he or she does not get the stated speed.
- The most interesting observation during this preliminary trial is that most users' two WTP values are pretty close to each other. The average difference between the first and second WTP values (among all valid replies) over

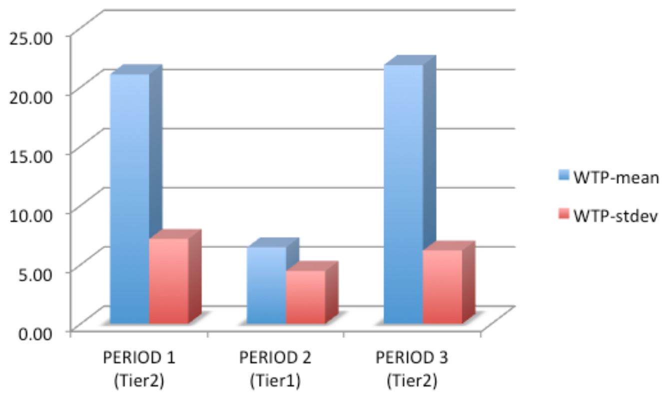

Fig. 3. WTP mean and stdev of the three periods among 12 trial participants

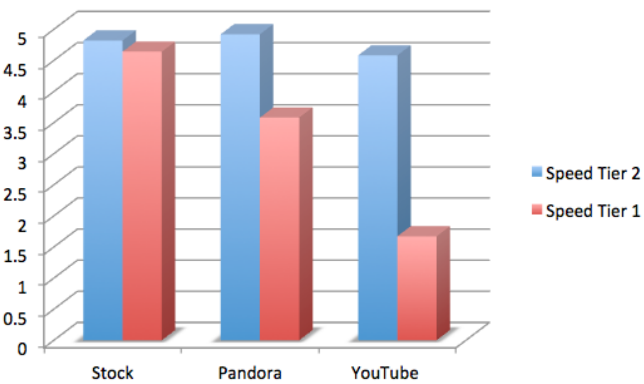

Fig. 4. Application rating on three apps: Stock, Pandora, and YouTube, based on speed tiers 1 and 2

the corresponding higher WTP is only 21.47%, while the speed difference ranges from 4 to 8 times between the two tiers.

The preliminary trial is not conclusive due to the limited number of samples, but the last point suggests that most subjects seem to have a limited dynamic range of valuation for the wireless services regardless of the speed tiers offered - and this points to opportunities for auction-based pricing, where a smartphone user would name a fixed price (based on their budget) and have their speed tier be ranked accordingly. We are currently expanding the trial to up to 1000 users in order to get enough statistical significance and to answer various pricing-related questions for service providers.

Second trial: Measuring WTP through application-rating using three treatments in the Lab

In the second setup, we invited 12 participants to come to a lab setup and conducted a 45-minute session (with short interviews) where the subject would go through three speed-tier treatments (tier2, tier1, tier2), where tier2 is 16 times the speed of tier1. Due to the lab setup, the actual speed the user experienced on the Wi-Fi network is close to the maximum throughput allowed by the speed tier. During each treatment, we also asked each subject to perform several tasks using selected applications (Stock, Pandora, YouTube, Weather, and Google Map) and rate the experience under the throttled speed. The reason for the third treatment with the same speed tier used during the first treatment was to see if the user would appreciate the higher speed tier more after experiencing the lower speed tier. We have the following observations after the trial:

- Unlike the first setup in the field, the stable speed due to the lab environment and the 16 times difference in speed between the two tiers has indeed created a major difference in the WTP values for the two speed tiers, as shown in Figure 3. The mean for WTP in period 1 (tier 2) was $21.08, while the mean for WTP for period 2 (tier 1) was only $6.48. There was a slight increase of WTP from period 1 to period 3 (same speed tier), which has a WTP mean of $21.88.
- Figure 4 shows the average rating of the application experience by the 12 participants under the two speed tiers on three apps with different levels of bandwidth requirements: Stock, a simple app with simple text/graphics, Pandora, a streaming music app, and YouTube, a video app. The users gave significantly lower ratings on YouTube (at the resolution of 360p) when using speed tier 1. The impact is less pronounced on Pandora, which buffers enough of the audio stream initially to provide a smooth listening experience even at a low speed tier. There was virtually no difference in the rating for the stock app.

Additional Lessons: We have also run into several unexpected problems during the construction of the testbed and we would like to briefly summarize our experience here:

- **VPN Performance:** It turns out that several older-generation smartphones could not achieve any throughput beyond 2Mbps (even with WiFi) when they use the VPN connection. This complicates our trial as we had to carefully map those phones only to speed tiers lower than 2Mbps.
- **Android Proxy-Redirection Issues:** While we are able to establish VPN connections to Android smartphones, the Cisco VPN appliance could not redirect them to specific proxy servers and ports that the certificates specify. We expect this issue to be resolved by Cisco and/or Android engineers in the future. In the mean time, we are exploring a cloud-based solution to bypass this issue.

IV. SpeedGate for Auction-Based Dynamic Pricing

SpeedGate can be used to support auction-based dynamic pricing algorithms through the addition of an auction manager. Figure 5 shows our design of the auction manager and its components. A bid for service-level adjustment is received through the bid interface manager and sent to the auction engine, which ranks each of the existing users connected to the testbed. User bids provide an indication of value each user places on receiving data services facilitated by the testbed. For example, users may spend more money in return for a tier level of service that provides higher speed tier, thereby resulting in faster data service performance. The auction engine works in conjunction with the resource manager to determine the number of available slots for different speed tiers, the user profile store to get all existing bids, and the user bandwidth monitor to guage the users' wireless channel quality, before determining the ranking of each user and allocating the speed tier through the bandwidth allocation engine.

As an example, the current SDP testbed has 4 proxies, each with the capability to support up to 64 VPN connections, resulting in a total of 256 VPN connections. The total bandwidth is constrained by the 64Mbps connection from the testbed to the Internet backbone. The core resource manager decides how to allocate the actual number of active VPN connections for each speed tier. For example, the testbed may allocate 16 16Mbps slots, 32 4Mbps slots, 80 1Mbps slots, and 128 256kbps slots based on the bandwidth constraints, typical resource utilization, the expected revenue, etc., among other factors. Note that not all active VPN connections would consume the maximum bit-rates allocated to them (and some may remain idle for a long time); therefore, the total bandwidth allocation may not exceed 64Mbps most of the times. The auction engine takes input from the bandwidth monitor and the user bids stored in the user profile store to rank all the current active user connections, and determines if the speed tier of any existing VPN connection should be adjusted (based on an auction algorithm similar to PFauc[2] in nature). If the speed of a connection needs to be adjusted, it sends the request to the bandwidth allocation engine, which sends an LDAP request to the certificate server (see Section II) to modify the proxy and port setting associated with the user's certificate and reset the VPN connection, which in turn causes the speed tier of the user to be adjusted. Note that our implementation is slightly different from the actual PFauc algorithm in that we adjust the speed tier at a "macro level", while PFauc makes low-level decisions every 2ms to determine the number of channelization codes to be allocated to each user.

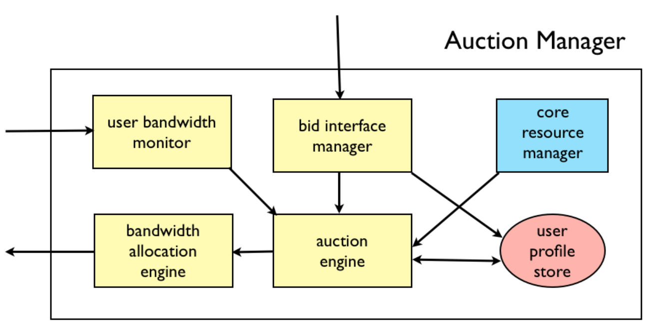

Fig. 5. Auction manager of SpeedGate

While many components of the testbed have been built, we are still at the design stage for the smartphone UI for GSP. In the illustrated example of Figure 6, the price indicator "$" on the top status bar is a monetary value to represent the estimated cost of joining a higher and/or lower tier of service. The core resource manager is tasked to measure the activity levels. Depending on the degree of congestion and the current bids from other users, a corresponding price indicator may be presented in the example status bar.

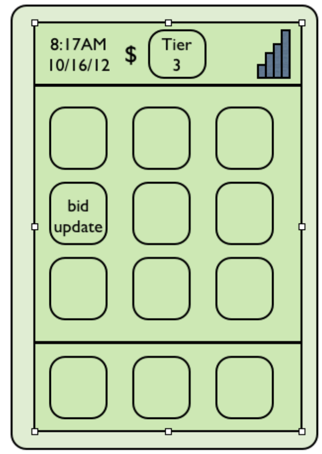

Fig. 6. Smartphone UI for Smart Data Pricing

The bandwidth tier indicator in the status bar informs a user of a type of bandwidth service enabled on the wireless device. For example, a first tier may represent a lowest level of priority (such as 256kbps) for data communications, while a fourth tier may represent a highest level of priority (such as 8Mbps). The bandwidth tier indicator may flash, change color and/or the wireless device may beep and/or vibrate when there is an opportunity to move to a different tier plan based on the user's preference. The user may invoke a bid via a button press. An example app "bid update" may correspond to a bidding application that allows the user to enter a bid value amount in an effort to increase a tier status, decrease a tier status and/or otherwise manage one or more bids. In the event that the entered bid amount is sufficient to raise a current tier status, the example bandwidth tier indicator in the status bar may change its appearance.

In the future, we would like to explore how the SDP testbed can be enhanced to support other well-known pricing schemes for side-by-side comparisons on resource utilization, social welfare, revenues/profits, etc.:

- Paris Metro Pricing[11]: The SDP testbed can simulate several logically separate channels by throttling at the proxy level (say at 16 Mbps) instead of at the port level. The proxy servers that host the channels with higher prices would attract less traffic and hence likely to provide better QoS since fewer users would share the 16 Mbps bandwidth.
- Time-Delayed Pricing (TDP)[8]: TDP addresses the congestion problem by considering when a user consumes data, in addition to how much is used. The TUBE system was used to conduct a trial based on TDP and it computes TDP prices so as to balance the cost of congestion during peak periods with that of offering lower prices in less congested periods. The SDP testbed can simulate TDP by removing all throttling through proxies, while adjusting prices charged to users by monitoring the VPN connection times through logs captured on the syslog server; however, many additional components from TUBE, such as their feedback-control mechanism

and their smartphone GUI must be added as well to facilitate a full-scale simulation.

V. Conclusion

In this paper, we presented SpeedGate, a smart data pricing testbed that allows a service provider to experiment with different dynamic pricing strategies on mobile data. SpeedGate maintains VPN connections to smartphones that can be adjusted to different speed tiers based on data pricing strategies. The testbed can maintain persistent VPN connections as smartphone users switch among different wireless networks (3G, 4G/LTE, and WiFi networks). We report preliminary results on two trials with a total of 29 users for assessing their willingness to pay (WTP) for various speed tiers. Preliminary observations point to the challenges of QoS guarantees through speed tiers, the limited dynamic range of WTP values from each user for different speed tiers, and potential opportunities for auction-based dynamic pricing. We also describe how the testbed can be extended with an auction engine to handle Generalized-Second-Price (GSP) auctions for dynamic pricing.

VI. Acknowledgments

The authors gratefully acknowledge the contributions of Karthik Kannan and Harsha Raman in in managing the first trial, and the administrative support from Steven Gao, Gang Yao, Robert Hooper, and Fu Che during various stages of the trials, and the support and guidance from Dana Tardelli on this project.

References

[1] Cisco anyconnect secure mobility client for mobile platforms data sheet. https://bitly.com/shorten/.

[2] Y. Chen, R. Jana, and K. Kannan. Using generalized second price auction for congestion pricing. In *Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE*, pages 1–6. IEEE, 2011.

[3] Cisco visual networking index: Global mobile data traffic forecast update, 2010-2015. Cisco White Paper, Feb. 2011. http://bit.ly/qjQmn.

[4] Cisco secure remote access cisco asa 5550 series ssl/ip sec vpn edition. http://bit.ly/6emG5I.

[5] D. D. Clark. Internet cost allocation and pricing. *Internet Economics*, 1997.

[6] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang. Pricing in computer networks: Motivation, formulation and example. In *IEEE/ACM Transactions on Networking*, 1993.

[7] Contingent valauation. http://en.wikipedia.org/wiki/Contingent_valuation.

[8] S. Ha, S. Sen, C. Joe-Wong, Y. Im, and M. Chiang. Tube: time-dependent pricing for mobile data. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 247–258. ACM, 2012.

[9] J. MacKie-Mason and H. Varian. Pricing the internet. In *In Public Access to the Internet - Prentice Hall*, 1995.

[10] A. Odlyzko. Paris metro pricing for the internet. In *Proc. of 1st ACM Electronic Commerce*, 1999.

[11] A. Odlyzko et al. Paris metro pricing for the internet. In *Electronic Commerce: Proceedings of the 1 st ACM conference on Electronic commerce*, volume 3, pages 140–147, 1999.

[12] 3gpp ts 23.203 - policy and charging control architecture. http://www.3gpp.org/ftp/Specs/html-info/23203.htm.

[13] L. Rizzo. Dummynet: a simple approach to the evaluation of network protocols. *ACM SIGCOMM Computer Communication Review*, 27(1):31–41, 1997.

[14] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. Lazar. Pricing, provisioning and peering: Dynamic markets for differentiated internet services and implications for network interconnections. In *IEEE Journal on Selected Areas in Communications*, 2000.

[15] S. Sen, C. Joe-Wong, S. Ha, and M. Chiang. Pricing data: A look at past proposals, current plans and future trends. In *arxiv.org/abs/1201.4197*, 2012.