Increasing Wireless Revenue with Service Differentiation

SIAMAK AYANI and JEAN WALRAND
Department of Electrical Engineering and Computer Sciences
University of California at Berkeley, Berkeley, CA 94720, USA

ABSTRACT

New technologies such as IEEE 802.16 (Wi-MAX) [2] and IEEE 802.11e (Wi-Fi with QoS) [1] enable differentiated services. In this study we explore the potential for increasing the revenue of Internet service providers with differentiated services.

For concreteness, we focus on two generic services that we call “voice” and “data.” However, the analysis applies to services that have different requirements such as video streaming and networked games. Using analytical models we examine the conditions under which the operator can improve his revenue with service differentiation. We verify the results using a simulation of a more complex model of the network.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

1. INTRODUCTION

For a given link technology, differentiated services can enable richer applications by reducing delay or increasing throughput. That fact led to the recent development of wireless MAC protocols with differentiated services such as 802.11e and 802.16. In this study, we explore the possibility of increasing the operator revenue by implementing differentiated services. Our approach is to compare the revenue of two systems: one without and the other with service differentiation. To make the comparison fair, we assume that the two systems have the same link capacity.

Differentiated services can be useful only for links that are likely to be a bottleneck, such as shared wireless access links (e.g., Wi-MAX and metropolitan Wi-Fi). For such a link, differentiated services can protect applications that require more throughput or smaller delays against more elastic applications. For example, it is intuitively clear that isolating voice from web traffic might improve the voice quality at a small cost for web users. This improved user experience can result in increased revenue for the operator. For links with a higher bit rate, the quality of voice calls may be adequate without service differentiation. However, other applications such as IPTV, video conference, or networked games might benefit from such differentiation.

In this paper we propose a model that captures the economic aspects of introducing a premium service. In the model, the operator reserves a fraction of the link capacity for the premium service. The interesting situation is when the corresponding reduction of the capacity of the basic service reduces the quality of that service. The question in that case is whether the revenue from the premium service compensates the loss of some of that from the basic service. The novel aspect of our model is that reducing the capacity of the basic service reduces the set of subscribers to the network and, accordingly, the set of potential users of the premium service. Thus, in contrast with standard models of service differentiation, reducing the quantity of basic service also reduces the demand for the premium service. Instead of considering an abstract formulation of this effect, we construct a model representative of the behavior of users of network services.

Other studies have explored the pricing of differentiated service, either using game theory (see [3], [4], [5] among others), or a duplicated network approach ([6]). This paper's contribution is to focus on a model of a shared Wi-Fi network.

In section 2 we construct the model of service differentiation. We analyze that model in section 3. In section 4, we examine the revenue improvements for concrete applications. Section 5 reviews some simulation results of a more detailed model of the system. The paper concludes with section 6 that reviews the conclusions of the study.

2. MODEL

The system without and with service differentiation is illustrated in Figure 1. The system without differentiation is a particular case of the system with differentiation with $\alpha = 1$. Consequently, it suffices to consider the system shown on the right of the figure. There is a pool of M actual subscribers that generate data connections and voice calls. The operator sets aside a fraction $1-\alpha$ of the capacity C to serve voice calls. The users first attempt to place a voice call using the basic service. The quality of the voice call is satisfactory as long as most voice packets experience a delay less than T_0. If a user is not satisfied with the quality of the voice call, he

Permission to make digital or hard copies of all or part 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 bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
QOSIM 2008, March 03, Marseille, France
Copyright © 2008 ICST 978-963-9799-20-2
DOI 10.4108/ICST.SIMUTOOLS2008.3068

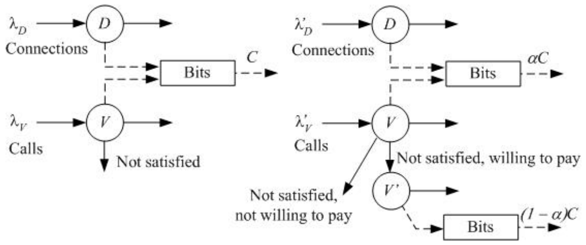

Figure 1: System without (left) and with (right) service differentiation.

can request a premium service if he is willing to pay an extra fee of p. In the extreme case, the operator could reserve a negligible capacity for the basic service. What prevents the operator from making that choice is that it would lose subscribers. The operator adapts α and p to the user demand to maximize the combined revenue of the basic and premium services. Thus, a central aspect of our model is to capture the loss of revenue incurred by reducing the capacity of the basic service.

We consider a pool of N potential subscribers to the access network. A user spends a fraction of his time using the network. In our simplified model, we consider only two activities: file transfers and voice calls. Each file transfer takes some time, both for the download and for using the file, such as reading a web page. We model that behavior by assuming that a user who is engaged in "data activity" transfers R_D bits per second, on average. For instance, if a user who is actively browsing the web downloads a 1MByte web page every minute, then $R_D = 8 \times 10^6$bits/60sec $\approx$ 133kbps. A typical user is engaged in a data activity only some fraction f_D of a busy period of the day. For instance, a user may browse the web ten minutes between 2:00 pm and 4:00 pm which may be the busy time that we should worry about. In that case, $f_D = 1/12$. Ideally, the network is designed so that it can meet the needs of the users during the busy periods. Thus, if the data capacity is αC, then one hopes that

$$Nf_DR_D \leq \alpha C.$$

In fact, one might want some safety margin to account for the randomness of the number of data connections. For simplicity, let us say that this margin is already included in R_D.

What if C is too small to meet these users' needs? We assume that the provisioning is similar for all the access points of the operator and that the number of subscribers to that operator's network adjusts itself based on the perceived quality of the network. That is, the number of subscribers falls if the network capacity is too small. One could argue that the operator would increase the capacity of the network to meet the potential demand. However, it is conceivable that an operator might choose to favor premium services, even at a loss of subscribers for the basic service. It is precisely that choice that we want to explore. We model that effect by having the actual number of subscribers equal to $M(\alpha)$ where

$$M(\alpha) := \min\{N, \alpha\frac{C}{f_DR_D}\}. \qquad (1)$$

That is, N is the potential pool of users. If α is too small to accommodate them, the network loses some of the users to a competitor. This model captures the cost of insufficient capacity of the basic service. We model the revenue of the basic service as $p_0M(\alpha)$ where p_0 is the subscription price for the basic service, say per user per month.

The $M(\alpha)$ network subscribers are engaged in voice calls some fraction f_V of a busy epoch. A user engaged in a voice call generates a bit rate equal to R_V bps. The voice users can place voice calls using the basic service. The quality is acceptable if there are not too many data connections. Otherwise, these users have the option of paying a price p per second and use a premium service. Let us assume that the probability that the voice quality is not acceptable is equal to $\rho(\alpha)$. A fraction $\phi(p)$ of the unsatisfied users are willing to pay p, so that they generate a bit rate $M(\alpha)f_V\rho(\alpha)\phi(p)R_V$ for the premium voice service and one hopes that

$$M(\alpha)f_V\rho(\alpha)\phi(p)R_V \leq (1-\alpha)C.$$

However, if $(1-\alpha)C$ is too small for this inequality to hold, then a number of potential premium voice calls are dropped. Thus, bit rate of voice calls that the premium service supports is about

$$\min\{M(\alpha)f_V\rho(\alpha)\phi(p)R_V, (1-\alpha)C\}.$$

Dividing that number by R_V gives the number of voice calls that the premium service supports at a typical busy time. Each voice call generates revenue equal to p per second. Multiplying that number by the number of busy seconds in a typical month gives the revenue per month. Equivalently, we can re-normalize p so that it is expressed in revenue per month, as p_0 is. The revenue of the premium service is then

$$p\min\{M(\alpha)f_V\rho(\alpha)\phi(p), (1-\alpha)\frac{C}{R_V}\}.$$

Summing up, the revenue of the operator is given by

$$R(\alpha, p) = p_0M(\alpha) + p\min\{M(\alpha)f_V\rho(\alpha)\phi(p), (1-\alpha)\frac{C}{R_V}\}.$$

Figure 2 illustrates that function, for different values of the parameters.

3. ANALYSIS

The following result explores when a premium service improves revenue.

Theorem 1. *Assume that $\rho(\alpha)$ and $\phi(p)$ are non-increasing functions of their argument. A premium service improves revenue if*

$$Nf_DR_D \leq C \text{ and } \rho(1) > 0 \qquad (2)$$

or if

$$Nf_DR_D > C \text{ and } \rho(1) > 0 \text{ and } \phi(p_0\frac{R_V}{f_DR_D}) > 0. \qquad (3)$$

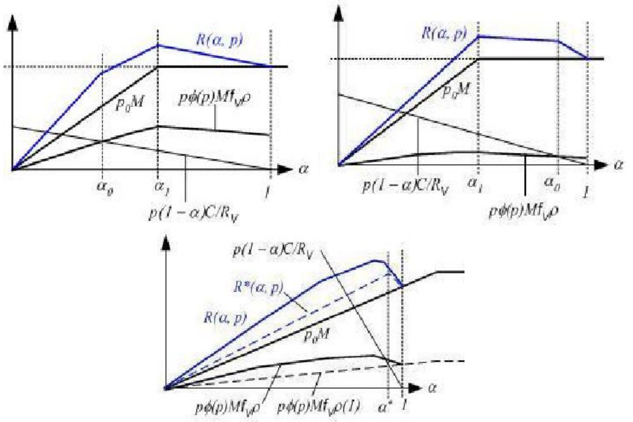

Figure 2: Revenue $R(\alpha, p)$ as a function of α for different parameters.

Condition (2) means that the capacity C is more than sufficient to accommodate all the data users. In that case, some of the capacity can be used to provide a premium service without reducing the revenue of the basic service. Under these assumptions, customers use the premium service if the basic service is not always acceptable, even though its capacity is large. This condition holds for MAC protocols where the quality of the voice calls may be unacceptable even though there is enough capacity to handle the data connections. This is the case for Wi-Fi networks where only a few data connections suffice to make the delay jitter too large for the voice quality to be acceptable. We believe that the situation is similar for the best effort service of Wi-MAX. Of course, if the link rate is very large, then this problem does not exist for voice, although it might for other applications such as games or video conferences.

Conditions (3) mean that the capacity C is not sufficient to accommodate all the data users. Using some of that capacity to provide a premium service reduces the revenue of the basic service. However, the additional revenue of the premium service is larger than the lost revenue from the basic service. The conditions are sufficient but not necessary.

The theorem shows that, assuming that the cost of providing the premium service is small, an operator that implements it has a competitive advantage. Indeed, with the same capacity, such an operator can generate a larger revenue with the same price of the basic service. Accordingly, that operator can lower the price of the basic service and remain profitable and gain an advantage over its competitors that do not introduce a premium service.

Proof of Theorem 1:

(a) Assume that condition (2) holds. That situation corresponds to the two upper diagrams of Figure 2. Let $\alpha = \alpha_1 := Nf_DR_D/C$. We find that

$$R(\alpha_1, p) \geq p_0N + p(1-\alpha_1)C/R_V > R(1, 0) = p_0N,$$

so that a premium service with any $p > 0$ increases revenue.

(b) Assume that conditions (3) hold. Note that

$$R(\alpha, p) \geq p_0M(\alpha) + p\min\{M(\alpha)f_V\rho(1)\phi(p), (1-\alpha)\frac{C}{R_V}\}$$

and designate the right-hand expression by $R^*(\alpha, p)$.

The lower part of Figure 2 shows the function $R^*(\alpha, p)$. Thus, it suffices to show that there is some choice of α and p so that $R^*(\alpha, p) > R^*(1, p) = R(1, 0)$. Consider the slope of $R^*(\alpha, p)$ at $\alpha = 1$. We find

$$\frac{d}{d\alpha}R^*(\alpha, p)|_{\alpha=1} = p_0\frac{C}{f_DR_D} - p\frac{C}{R_V} \text{ if } \alpha^* < 1 \qquad (4)$$

where α^* is the value of α where $p\phi(p)M(\alpha)f_V\rho(1) = p(1-\alpha)C/R_V$. That is, after some algebra,

$$\alpha^* = \frac{f_DR_D}{f_DR_D + R_V\phi(p)}. \qquad (5)$$

Identity (4) shows that the slope of R^* at $\alpha = 1$ is negative if

$$p > p_0\frac{R_V}{f_DR_D}.$$

If conditions (3) hold, there is some p that satisfies the condition above so that $\phi(p) > 0$ and, consequently, $\alpha^* < 1$.

□

The maximum revenue can be determined if one knows the function $\rho(\alpha)$. We examine a simple case in the fact below that one can verify using simple algebra.

COROLLARY 1. *Assume that the maximum revenue occurs for some α such that $\rho(\alpha) \approx \rho(1)$. Then*

$$\max_\alpha R(\alpha, p) \approx C\frac{p_0 + p\phi(p)\rho(1)f_V}{f_DR_D + \phi(p)\rho(1)f_VR_V}.$$

One can then find the price p that maximizes the revenue by maximizing that expression.

4. APPLICATIONS

In this section we explore a number of concrete applications and study the revenue improvements with service differentiation.

Consider an 802.11g network. We choose, as explained in section 2, $f_D = 0.1, f_V = 0.05, R_D = 150$kbps, $R_V = 35$kbps, $C = 24$Mbps, and $\phi(p) = \exp\{-p/\theta\}$. To estimate ρ, we argue as follows. Assume that there are N users that transfer files about once every 10 seconds when they are active and that the files take about 5 seconds to be transferred. The file transfers correspond to a Poisson process with arrival rate $\lambda = N \times f_D/10$. The mean service time is $\mu = 1/5$. According to experiments, a voice call is unacceptable as soon as one file is being transferred. Say that a user decides to switch to the premium service if the voice call is unacceptable within the first 20 seconds. The probability that no file is being transferred during $T = 20$ seconds is approximately $1 - \rho = \exp\{-\lambda/\mu\}\exp\{-\lambda T\}$. For instance, if $N = 10$, one finds $\lambda = 0.1$ and $\rho \approx 0.87$. Since there a lot of uncertainty about the activity of users, let us choose a conservative estimate $\rho = 0.2$. For these values, conditions (2) hold.

For video streaming, $\rho(\alpha)$ is the probability that the available capacity becomes less than 250kbps (e.g.) during one video transfer. This is the probability that $D + V$ reaches a value k such that $\alpha C/k < 250$kbps.

For networked games, one can expect $\rho \approx 1$. Indeed, without service differentiation, the delay of game packets is likely to be almost always excessive. In this case, conditions (3) are likely to hold for most systems.

5. SIMULATIONS AND RESULTS

It is desirable to study the revenue without detailed knowledge about $\phi(p)$ and $\rho(\alpha)$. For this purpose, we developed a discrete event simulation platform to investigate the impact of various values of p and α on the revenue. The simulation platform can be used for a range of applications, such as voice and video streaming mentioned in section 4. Below we describe the simulation of a system that provides services for data transmission and VoIP.

5.1 Simulation of data and voice services

Here we consider a case where users are connected to a local wireless access point of an ISP that provides a basic voice service for its subscribers who pay a subscription fee of p_0 per month. A subscriber can chose a premium voice service by paying an additional fee of p per minute. We investigate the impact of pricing on the generated traffic and the revenue of the ISP.

The system allocates a fraction $1 - \alpha$ of its capacity C to serve the premium voice calls. Further, we assume that data connections and voice calls arrive as Poisson processes with rates λ_D and λ_V respectively. We denote the number of voice users in the basic and premium service at a given time as V and V' respectively. We assume that the duration of a voice connection is exponentially distributed with mean $1/\mu_V$. Similarly, we denote the number of subscribers using data service by D. We assume that the data connections correspond to file transfers that have some exponentially distributed length with mean L_D.

The voice users experience an unacceptable delay (e.g., $T > T_{threshold}$) if the basic service is congested. Such customers would leave the system if it did not provide any other voice service. However, in the two-class differentiated service system, these users have the option of joining the premium service. They do so if they are willing to pay the extra cost for using the premium service, otherwise they leave the system. One user's willingness to pay is independent of that of the other users and has some probability $\phi(p)$ of exceeding p. In order to guarantee a better QoS for the premium service, a limitation is set on the number of customers that can use this service simultaneously. Furthermore, we assume that the users make a decision to join the basic or the premium service upon arrival depending on number of customers already using the basic and the premium service. In order to capture the cost of allocating insufficient capacity to the basic service and hence shrinking the pool of actual users, we assume that $\lambda_D = \min\{\alpha\lambda_{D_0}, \lambda_{D_{max}}\}$ and $\lambda_V = \min\{\alpha\lambda_{V_0}, \lambda_{V_{max}}\}$. Using the above defined rates we model the system as a continuous-time Markov chain and conduct a discrete event simulation of the system.

5.2 Simulation results

Figure 3 shows the obtained revenue for different values of α when the price of the premium services varies from 1 to 10. The figures also illustrate that, for the chosen system specifications, an introduction of a premium service improves the revenue.

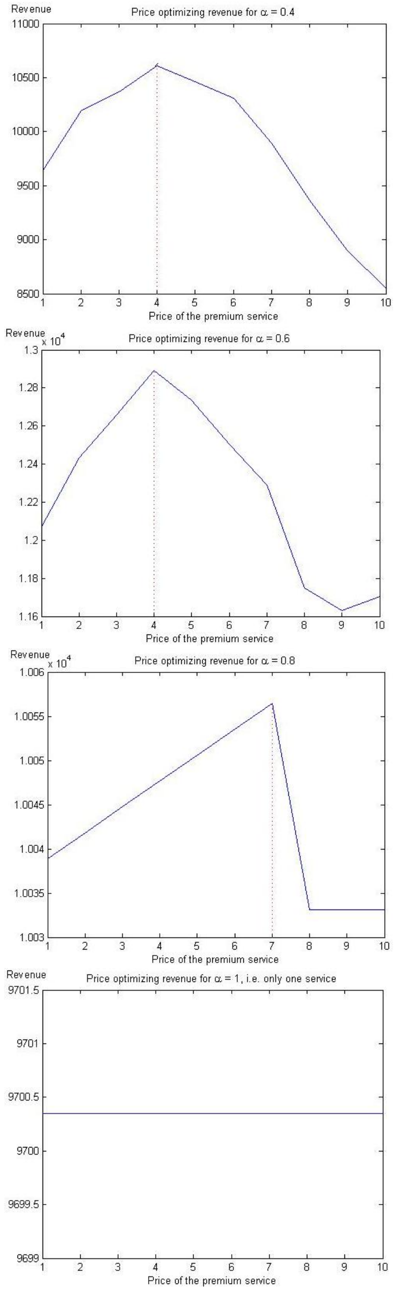

Figure 3: Obtained revenue for various α and a range of prices, p, charged for the premium service. The bottom figure shows the revenue for the single service scenario.

5.3 Finding the optimal price

Figure 3 shows that the revenue varies with α and p and has a local maximum. Below, we present an iterative algorithm that, for a given α, finds the optimal value p, i.e., the value that maximizes the revenue. The algorithm increases the price by E at each step as long the revenue is increasing. However, if the revenue decreases then we rollback and increase the price only by $E/2$ (i.e. we set $E = E/2$). If the new revenue is increasing then we keep that value, otherwise we roll back to the previous one.

```
1. Set $p = 0$ and run the simulator to calculate revenue $R_o$
2. Set $R_{opt} = R_o$
3. Set $p = p + E$ and run the simulator to calculate revenue $R_{new}$
4.   • If $R_{new} > R_{opt}$
         – $R_{opt} = R_{new}$
         – go to 3
     • elseif $E > \epsilon$
         – $E = E/2$
         – go to 3
5. end
```

Figure 4 illustrates the result of running the above algorithm in the above mentioned simulation platform. We can now see that the $\alpha_{optimum} = 0.6$, $p_{optimum} = 4$ and that maximum revenue will be about 32% greater than the revenue of the single service system.

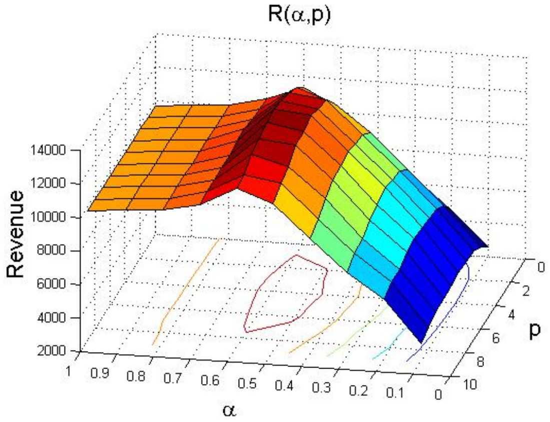

Figure 4: The revenue for a range of α and p, where $\phi(p)$ has a negative exponential distribution with expected value 5.

6. CONCLUSIONS

In this paper we have investigated the conditions under which ISPs can improve their revenue by introducing differentiated services. We have introduced a model that captures the cost of allocating insufficient capacity to the basic service, and studied a generic case where two services that provide different QoS are available to customers. We have shown the conditions under which service differentiation actually leads to increased revenue. In order to elaborate on these issues, we have chosen to do a simulation of a typical system and shown how the revenue varies with the price charged for the premium service and with the capacity provisioning. We also suggest a method that can be used by operators to find the optimum price and system capacity allocation.

The analysis of the previous section assumes that the operator knows the key parameters of the system, such as the functions $\phi(p)$ and $\rho(\alpha)$. In practice, that information is difficult to obtain. Hence, it is useful to derive mechanisms to adjust p and α without requiring a detailed knowledge of the system characteristics. We plan to explore adaptive pricing mechanisms based on stochastic approximations to discover the optimal price and provisioning.

Another aspect of the system that we plan to explore is the heterogenous nature of customers. Some users value premium services more than others and have a different willingness to pay for such services. The question is whether an operator should specialize in premium services to capture those customers.

7. REFERENCES

[1] IEEE802.11 WG, “IEEE802.11 Standard,” http://www.ieee802.org/11/

[2] IEEE802.16 WG, “IEEE802.16 Standard,” http://standards.ieee.org/getieee802/802.16.html

[3] L. He and J. Walrand, “Pricing Differentiated Internet Services,” *INFOCOM 2005.*

[4] P. Marbach, “Pricing Differentiated Services Networks: Bursty Traffic,” *INFOCOM 2007.*

[5] J. Musacchio and J. Walrand, “Game Theoretic Modeling of WiFi Pricing,” *Allerton 2003.*

[6] A. Odlyzko, “Paris Metro Pricing: The minimalist differentiated services solution,” *ACM Conference on Electronic Commerce.* 1999.