

Time-Dependent Network Pricing and Bandwidth Trading

Libin Jiang*, Shyam Parekh*[†], Jean Walrand*

*Dept. Electrical Engineering & Computer Science, University of California, Berkeley

[†]Bell Laboratories, Alcatel-Lucent, Murray Hill, NJ 07974

Email: {ljiang,shyam,wlr}@eecs.berkeley.edu

Abstract—The usage of a network usually differs significantly at different times of a day, due to users’ time-preference. This phenomenon is also prominent in the market of “bandwidth-on-demand”, since the demand is typically higher during large events. Thus, an unselfish “social planner” should deploy a proper pricing scheme to reduce congestions and achieve efficient use of the network (i.e., maximize the “social welfare”); whereas a selfish service provider (SP) can exploit the time-preference to increase its revenue. In this paper, we present a model to study the important role of time-preference in network pricing. In this model, each user chooses his access time based on his preference, the congestion level, and the price he would be charged. Without pricing, the “price of anarchy” (POA) can be arbitrarily bad. We then derive a simple pricing scheme to maximize the social welfare. Next, from the SP’s viewpoint, we consider the revenue-maximizing pricing strategy and its effect on the social welfare. We show that if the SP can differentiate its prices over different users and times, the maximal revenue can be achieved, as well as the maximal social welfare. However, if the SP has insufficient information about the users and can only differentiate its prices over the access times, then the resulting social welfare can be much less than the optimum, especially when there are many low-utility users. Otherwise, the difference is bounded and less significant.

I. INTRODUCTION

The usage of a communication network usually differs at different times of a day, because each user has its own preference over the access time, and the distribution of time-preference among a population affects the distribution of actual traffic. Also, the technological advances in networking has made it possible for “on demand” provisioning of bandwidth for applications that require short-term bandwidth assignments, such as in the event of sporting or cultural activities. In this market, the time-preferences of the bandwidth buyers (or “users”) are especially prominent since such events draw much higher demand than usual. Also, the users’ different levels of interests on different events lead to different preferences over the time they wish to obtain the bandwidth.

Without proper use of pricing, traffic delays caused by congestion during busy times can be unacceptably high, leading to inefficient utilization of the network. Intuitively, by charging higher prices during busy times, excessive congestion can be avoided, and the revenue collected by the service provider (SP) can also be increased.

This work is supported by the National Science Foundation under Grant NeTS-FIND 0627161. Sincere thanks to the collaboration of Bell Labs.

A complication in this scenario is that there is no simple “demand function”, which relates the price to the user demand. The demand of the users are different at different time, but one cannot use separate demand functions for different times, because charging a higher price at a specific time will not only reduce the demand at this time, but also drive the users to migrate to other times. Therefore the demands at different times are closely correlated.

In this paper, we propose a simple model to study how time-preference affects the game between the users and the service provider. Specifically, it is modeled as a *selfish routing problem with user preference*. We derive the pricing scheme that achieves efficient use of the network (i.e., maximizes the “social welfare”, from a social planner’s standpoint), and the one that maximize the total revenue (from a SP’s standpoint). We will show later that the resulting social welfare under the two schemes may or may not be the same, depending on how much user information the SP has.

Selfish routing has been studied in various contexts. A classical model is by Wardrop [3] for the transportation network. The concept of “Price of Anarchy” (POA) was introduced in [4], measuring how bad the worst-case Nash Equilibrium (NE) is compared to Social Optimum (“SO”), in terms of some objective function (“social welfare” or “social cost”). Roughgarden et al showed that for selfish routing without pricing, and with the “social cost” is defined as the total delays of all users, the POA can be unbounded; but if the delay function is affine, the POA is upper-bounded by $4/3$ [2]. A number of other works focus on networks with parallel “links”. These links are owned by a SP, or separately owned by a number of SP’s. The SP’s seek to maximize their own revenue, and are part of the game in addition to the users. Reference [5] shows that if there is one SP (monopolist), social optimum can be achieved at NE (where the SP’s revenue is maximized and the users reach Wardrop Equilibrium), if there are a large number of users and each user has a fixed amount of data to send/receive. If there are multiple SP’s (oligopolists) to compete for the users, [6] and [7] derived the POA’s with inelastic traffic and elastic traffic. In [8], Musacchio et al derived the same POA using a different approach for elastic traffic. Our model here involves parallel links (analogous to “time slots”, as detailed later), a monopolist, and assumes each user has a fixed amount of traffic (for simplicity). But

different from all the works above, each user has preference over different time slots. In other words, a user would achieve different utilities if he would choose different time slots to access the network.

Reference [9] studied congestion-dependent pricing using a dynamic programming formulation. Interestingly, some conclusions of [9] are similar to ours qualitatively, although the approach used here seems more tractable.

II. BASIC MODEL AND NASH EQUILIBRIUM

A. The Model and its Nash Equilibrium

Divide the time in a certain period (for example, one day in the case of “time-of-day pricing”; or one week in the case of “bandwidth-on-demand” trading) into N equal time slots. For simplicity, assume that each user picks one time slot to access the network, and has a fixed amount of traffic. (In the conclusion, we will discuss the extension where the users access multiple slots, and use different types of network service, such as data or voice.) Assume there are K classes of users, each of which has a profile of time-preferences. (For example, in the case of “time-of-day pricing”, office workers and college students can be defined as two different classes. In the case of “bandwidth-on-demand” trading, sports broadcasters and music broadcasters are two different classes.) The number of class- k users is $d_k = \sum_{i=1}^N d_k^i$, where d_k^i is the number of class- k users choosing slot i . Let u_k^i be the “utility” of a class- k user if he chooses slot i , reflecting his preference. Meanwhile, the user experiences congestion delay, depending on the number of users choosing the same time slot. Then, the “payoff” for a class- k user who picks slot i is

$$v_k^i = u_k^i - l_i(x_i)$$

where $l_i(\cdot)$ is the congestion delay function, which is assumed to be increasing and convex; and $x_i := \sum_{j=1}^K d_j^i$. If a particular time slot is preferred by many users, leading to high congestion, then the individual payoffs may not be high for those who choose the slot.

For some users, it is possible that no matter which time slot they choose, their payoffs are negative due to large congestion delays or low utilities. If they choose not to use the network, their payoff is 0. Thus, these users will not use the network.

Note that the time slots here are analogous to the parallel links in the network model ([5], [8]), since in both cases the users strategically choose from a number of options (time slots or links). But in our problem, unlike previous works, a user would get different utilities in different times due to his time preferences, a phenomenon important in network pricing and bandwidth trading.

Now assume that the service provider or a social planner charges a price p_k^i to each class- k user in slot i (If $p_k^i = 0, \forall k, i$, then it is reduced to the case of no pricing.), then the payoff of this user becomes

$$f_k^i := u_k^i - l_i(x_i) - p_k^i \quad (1)$$

and each user tries to maximize its payoff f_k^i in the Nash game.

We have the following about the existence of NE.

Proposition 1: Given any fixed price vector $\mathbf{p} := \{p_k^i, k = 1, 2, \dots, K, i = 1, 2, \dots, N\}$, there exists a pure-strategy Nash equilibrium.

Proof: At Nash equilibrium, the following convex optimization problem is solved:

$$\begin{aligned} \max_{\mathbf{d}} \quad & \sum_{i=1}^N \sum_{k=1}^K (u_k^i - p_k^i) d_k^i - \sum_{i=1}^N \int_0^{x_i} l_i(y) dy \\ \text{st} \quad & \sum_{i=1}^N d_k^i \leq d_k, \forall k; d_k^i \geq 0, \forall k, i \end{aligned} \quad (2)$$

where $x_i = \sum_{k=1}^K d_k^i$ is the total number of users in slot i .

Indeed, the KKT condition [11] for (2) is that for each k , there exists a λ_k , such that

$$\begin{cases} u_k^i - p_k^i - l_i(x_i) = \lambda_k & \forall d_k^i > 0 \\ u_k^i - p_k^i - l_i(x_i) \leq \lambda_k & \forall d_k^i = 0 \end{cases} \quad (3)$$

This is exactly the condition for a NE. ■

The “social welfare” $V(\mathbf{d})$ is defined as the sum of the payoffs of the users and the SP (or the social planner), given an allocation $\mathbf{d} := \{d_k^i, k = 1, 2, \dots, K, i = 1, 2, \dots, N\}$:

$$\begin{aligned} V(\mathbf{d}) &:= \sum_{i=1}^N \sum_{k=1}^K f_k^i \cdot d_k^i + \sum_{i=1}^N \sum_{k=1}^K p_k^i d_k^i \\ &= \sum_{i=1}^N \sum_{k=1}^K [u_k^i - l_i(x_i)] \cdot d_k^i \end{aligned} \quad (4)$$

where $\sum_{i=1}^N \sum_{k=1}^K p_k^i d_k^i$ is the revenue collected by the SP or the social planner, and note that the term is canceled in $V(\mathbf{d})$. So a high value of social welfare means a good tradeoff between realizing the users’ utilities and reducing the congestions.

And the “Price of Anarchy” (POA) is the ratio between the social welfare at NE ($V_{NE} := V(\mathbf{d}_{NE})$) and at Social Optimum ($V_{SO} := V(\mathbf{d}^*)$, where \mathbf{d}^* is the maximizer of $V(\mathbf{d})$):

$$\rho := V_{NE}/V_{SO} \quad (5)$$

B. Price of Anarchy without Pricing

Without pricing, i.e., the price vector $\mathbf{p} = \mathbf{0}$, a simple example can show that the POA is arbitrarily bad. Let there be a single class, with population $d_1 = 1$. There is a single time slot with $u^1 = 1$. And assume that the delay function $l(x)$ is continuous, strictly increasing and satisfies $l(0) = 0$ and $l(1) = 1$.

At the NE, all users will use the time slot, in which case the payoff of each user is $u^1 - l(d_1) = 0$, the same as the payoff of not using the network. Therefore the social welfare is $V_{NE} = 0$.

At Social Optimum (SO), the number of users in the slot, x_{SO} , should be smaller than 1, in which case the social welfare is $V_{SO} = x_{SO}(1 - l(x_{SO})) > 0$. So $\rho = V_{NE}/V_{SO} = 0$.

Note that this example is not contrived. Similar situation is easy to happen in practice.

III. USING PRICING TO ACHIEVE SOCIAL OPTIMUM

With the price vector \mathbf{p} , (3) is satisfied at NE.

At the SO point, the KKT condition is that for each class k , there exists a β_k , such that

$$\begin{cases} \frac{\partial V(\mathbf{d})}{\partial d_k^i} = u_k^i - l_i(x_i) - x_i \cdot l'_i(x_i) = \beta_k & \forall d_k^i > 0 \\ \frac{\partial V(\mathbf{d})}{\partial d_k^i} = u_k^i - l_i(x_i) - x_i \cdot l'_i(x_i) \leq \beta_k & \forall d_k^i = 0 \end{cases} \quad (6)$$

(6) and (3) are identical if we let

$$p_k^i = x_i \cdot l'_i(x_i) \quad (7)$$

In other words, if the social planner sets the price as in (7), then the social optimum can be reached.

Interestingly, the price does not depend on the class k and the utility u_k^i , but only depends on the total number of users in slot i and the function $l_i(\cdot)$, which makes it easy to implement since the number of users is easy to observe. Essentially, the price can be interpreted the “externality” caused by a class- k user in slot i to other users—the existence of this user increases the delay of every other user in slot i by $l'_i(x_i)$, while the u_k^i component of every other user is not affected. (Similar forms of pricing also appear in some other scenarios [3], [5].)

IV. THE DISCRETE CASE

In the above we have assumed that there are many users, such that a single user is relatively “small”. As a result, the population can be modeled as a continuous variable. To understand the scenario where the sizes of users are “large”, and each user can only choose one time slot, we give a discrete formulation in the following. We will show that even if each user is priced similarly to the continuous case above, the social optimum may not be reached at some NE.

Say that there are M users. Each user chooses a time slot t in $\{1, 2, \dots, N\}$ to access the network. (We intentionally use different notations from the continuous case.)

User i chooses the time t_i to maximize

$$u_i(t) - l_t(x(t)) - p_t(x(t))$$

where

$$x(t) = \sum_{j=1}^M 1\{t_j = t\}$$

and

$$p_t(x(t)) = [x(t) - 1][l_t(x(t)) - l_t(x(t) - 1)].$$

Here, $u_i(t)$ reflects user i 's preferences. And the price $p(x(t))$ reflects the externality of user i if he chooses slot t .

The social welfare is $V = \sum_i u_i(t_i) - \sum_t x(t)l_t(x(t))$.

We show that every NE is “locally” socially optimal. That is, if user i switches from its slot at NE, $t_i = t$, to t' , the

social welfare cannot be increased. Indeed, if he makes the switch, then V is increase by

$$\begin{aligned} & u_i(t') - l_{t'}(x(t') + 1) - u_i(t) + l_t(x(t)) \\ & - x(t')[l_{t'}(x(t') + 1) - l_{t'}(x(t'))] \\ & - [x(t) - 1][l_t(x(t) - 1) - l_t(x(t))] \\ = & \{u_i(t') - l_{t'}(x(t') + 1) - \\ & x(t')[l_{t'}(x(t') + 1) - l_{t'}(x(t'))] \\ & - \{u_i(t) - l_t(x(t)) - \\ & [x(t) - 1][l_t(x(t)) - l_t(x(t) - 1)]\} \\ = & \{u_i(t') - l_{t'}(x(t') + 1) - p_{t'}(x(t') + 1)\} \\ & - \{u_i(t) - l_t(x(t)) - p_t(x(t))\} \\ \leq & 0 \end{aligned}$$

where the last step follows from the fact that user i selfishly prefers slot t to slot t' .

Although the NE is locally socially optimal, it may not be globally optimal. For example, assume there are only 2 users. User 1 is in class 1, user 2 is in class 2. User 1 prefer slot 1, user 2 prefer slot 2. Let $u_1^1 = 11, u_1^2 = 10; u_2^1 = 10, u_2^2 = 11$. Let the delay function in both slots be $l(\cdot)$, and $l(0) = 0, l(1) = 1, l(2) = 5$. Assume now user 1 is in slot 2, while user 2 is in slot 1. Then, no one will move unilaterally. But the social optimum is achieved when they exchange the slots.

The reason for the sub-optimality is the loss of convexity in this discrete model, such that the local optimum does not imply global optimum. If the user is allowed to split their traffic to more than one time slots (a reasonable assumption), then this problem will not occur even if the “size” of each user is large.

V. REVENUE MAXIMIZATION

Now we consider the pricing strategy of a revenue-maximizing SP, and how this affects the social welfare and its revenue.

A. SP has full information: Differentiated prices over time and classes

Under the idealized assumption that the SP has full knowledge of the users' classes and utilities u_k^i , he can price the users differently based on their classes and the slots they choose. Given an allocation $\mathbf{d} = \{d_k^i, k = 1, 2, \dots, K, i = 0, 1, \dots, N\}$, then for a class- k users in slot i , the SP can charge up to $u_k^i - l_i(x_i)$, making the user indifferent of using any slot or not using the network. Clearly, the revenue he can collect is the same as the social welfare. Therefore, when the social optimum is reached, the revenue of the SP is also maximized.

In this game, the monopolist SP has complete advantages since he can extract all the values of the users based on the user information he has, whereas the payoffs of the users are zeros. Therefore in this case, the SP's revenue is the highest among all possible pricing schemes. This is because in any other pricing scheme, (1) the social welfare must be smaller than the maximal social welfare (also the SP's revenue) achieved




Fig. 1. Loss of social welfare, when the SP maximizes its revenue with differentiated prices over time only

here, and (2) SP's revenue can not be larger than the social welfare (since part of the social welfare goes to the users).

B. SP has partial information: Differentiated prices over time

Usually, the SP does not have full information about the users' utilities. As a result, he may only differentiate its price over time. Unlike the previous case, maximizing revenue here is no longer the same as achieving social optimum. Let V_{MR} be the social welfare when the SP's revenue is maximized (where "MR" means "Maximal Revenue". It turns out that the ratio V_{MR}/V_{SO} can be arbitrarily close to 0, in sharp contrast to the previous case.

Consider the following example. There is only one time slot, "slot 1", along with slot 0. And the delay function $l^1(\cdot) = 0$. There is an infinite number of classes, class 1, 2, ..., each with a population $d_i = 1000$. Assume $u_1^1 = 3, u_i^1 = \frac{2}{i}, \forall i > 1$. This distribution of user utilities is shown in Fig 1, where x is the cumulative number of users and u is the utility level. The curve in Fig 1 means that there are 1000 users with utilities not less than 3; 2000 users with utilities not less than 1; 3000 users with utilities not less than $\frac{2}{3}$, etc. Since the delay function is assumed to be 0, Fig 1 can be viewed as a "demand function": for example, if the SP charges a price of $\frac{2}{3} - \epsilon$ (where $\epsilon > 0$ is very small), 3000 users will use the network since they will get positive payoff by doing that. (Fig 1 is also the inverse function of the complementary cumulative distribution function (c.c.d.f.) of the user utilities.)

Then, if the SP charges a price $p^1 = u_i^1 - \epsilon$, all users in class 1, 2, ..., i will go to slot 1. Then the revenue has many local maxima $R_i = u_i^1 \cdot 1000 \cdot i, \forall i = 1, 2, \dots$. Note that $R_1 = 3000, R_i = 2000 \forall i > 1$, so the price the SP charges is $p^1 = u_1^1 - \epsilon \approx 3$, and $V_{MR} \approx u_1^1 * d_1 = 3000$. However, optimal social welfare is achieved when all users are in slot 1, making $V_{SO} = 3000 + 2000 \sum_{i=2}^{\infty} \frac{1}{i} \rightarrow \infty$. Therefore, the ratio V_{MR}/V_{SO} can be arbitrarily close to 0.

C. Differentiated prices over time: A Continuum Model and the Price of Anarchy

The low value of social welfare (under the SP's revenue-maximizing pricing) in the last example clearly results from the specific distribution of the user utilities. To gain more understanding of this issue, we propose a continuum model in the following, which views the whole population as a single user with a certain utility function. Then, we derive a bound on the POA (i.e., efficiency loss) under some conditions on the utility function (or equivalent, the utility distribution of the population).

1) *Continuum population model*: Since there is only one price for each time slot, (2) can be written as

$$\begin{aligned} \max_{\mathbf{d}} \quad & \sum_{i=1}^N \sum_{k=1}^K u_k^i d_k^i - \sum_{i=1}^N p^i x_i - \sum_{i=1}^N \int_0^{x_i} l_i(y) dy \\ \text{st} \quad & \sum_{i=1}^N d_k^i \leq d_k, \forall k; d_k^i \geq 0, \forall k, i \end{aligned} \quad (8)$$

Note that given $x_i, i = 1, 2, \dots, N$, the last two terms of the objective function are fixed. Using a dynamic programming argument, if we define

$$\begin{aligned} U(x_1, x_2, \dots, x_N) & := \\ \max_{\mathbf{d}} \quad & \sum_{i=1}^N \sum_{k=1}^K u_k^i d_k^i \\ \text{st} \quad & \sum_{i=1}^N d_k^i \leq d_k, \forall k; d_k^i \geq 0, \forall k, i \\ & \sum_{k=1}^K d_k^i \leq x_i, \forall i \end{aligned} \quad (9)$$

then (8) is equivalent to

$$\begin{aligned} \max_{x_i, i=1, \dots, N} \quad & U(x_1, x_2, \dots, x_N) - \\ & \sum_{i=1}^N p^i x_i - \sum_{i=1}^N \int_0^{x_i} l_i(y) dy \\ \text{st} \quad & \sum_{i=1}^N x_i \leq D, x_i \geq 0, \forall i \end{aligned} \quad (10)$$

where $D := \sum_{k=1}^K d_k$.

$U(x_1, x_2, \dots, x_N)$ is the maximal total utility of users given $x_i, i = 1, \dots, N$. Note that it is concave according to the theory of convex optimization [11].

Under our previous assumption about classes, U is not smooth (not differentiable at some points). But if we consider a continuum of users with infinite population (i.e, the constraint $\sum_{i=1}^N x_i \leq D$ is dropped), and if the utility vector $\mathbf{u} = (u_1, u_2, \dots, u_N)$ of the users follow some continuous distribution $f(\mathbf{u})$, then it's reasonable to model U as a increasing, smooth, strictly-concave function.

We now quantify the relationship between the marginal utility $\nabla U(\mathbf{x})$ and the utility distribution $f(\mathbf{u})$. It is not required that the total population $\int f(\mathbf{u}) d\mathbf{u}$ is equal to 1.

For convenience, we assume that $f(\mathbf{u}) > 0, \forall \mathbf{u}$; and there is no atom in the distribution (i.e., there exists no \mathbf{u}' such that the utility vectors of a positive population of users are exactly equal to \mathbf{u}'). Let \bar{u}_i be the marginal utility in slot i (the smallest utility among the users in slot i), and $\bar{\mathbf{u}}$ be the vector of marginal utilities. Note that $\bar{\mathbf{u}} = \nabla U(\mathbf{x})$ is also the vector of dual variables of problem (9) given \mathbf{x} .

Since there is no atom in the distribution $f(\mathbf{u})$, one can imagine that there is an infinite number of ‘‘classes’’ in the population, and each class is very small and has a unique utility vector \mathbf{u} . Since we are solving a convex optimization problem, each ‘‘class’’ should choose one slot, denoted as i^* , that gives it the maximal ‘‘net utility’’, i.e., $i^* = \arg \max_i (u_i - \bar{u}_i)$, where \bar{u}_i is also the dual variable at slot i , as mentioned above.

Therefore, any user in slot i must have the following properties (Denote its utility vector as \mathbf{u}): (a) $u_i \geq \bar{u}_i$, (b) $u_j - \bar{u}_j \leq u_i - \bar{u}_i, \forall j \neq i$. So, the number of users in slot i satisfies

$$x_i = \int_{\bar{u}_i}^{\infty} f_i(u_i) Pr\{u_j \leq u_i - \bar{u}_i + \bar{u}_j, \forall j \neq i | u_i\} du_i \quad (11)$$

where $f_i(\cdot)$ is the marginal distribution function of the i 'th element of the population's utility vectors.

Equation (11) maps the vector $\bar{\mathbf{u}}$ to \mathbf{x} . Denote the mapping as $\mathbf{x} = G(\bar{\mathbf{u}})$. Then,

$$\nabla U(\mathbf{x}) = \bar{\mathbf{u}} = G^{-1}(\mathbf{x})$$

which determines $\nabla U(\mathbf{x})$ from the distribution $f(\mathbf{u})$. ($G(\cdot)$ is invertible since there is no ‘‘gap’’ in the distribution: $f(\mathbf{u}) > 0, \forall \mathbf{u}$.)

When there is only one slot, from (11), $x = \int_{\bar{u}}^{\infty} f(u) du := G(u)$; then $U'(x) = G^{-1}(x)$. For example, the curve in Fig 1 is actually $U'(x)$. But given a general distribution $f(\mathbf{u})$ with many slots, computing $\nabla U(\mathbf{x})$ may not be easy. In Section V-C3, however, we will give an example where it can be found explicitly.

2) *Price of Anarchy*: We now consider the price of anarchy in the game (where the SP is also a player). For convenience, let's denote $\mathbf{x} := (x_1, x_2, \dots, x_N)$.

At social optimum (SO), we maximize $V(\mathbf{x}) := U(\mathbf{x}) - \sum_{i=1}^N x_i \cdot l_i(x_i)$ subject to $\mathbf{x} \geq \mathbf{0}$. Denote the social optimum solution as $\mathbf{x}_S \geq \mathbf{0}$. Then, the optimality condition is that for some $\lambda \geq \mathbf{0}$,

$$\begin{aligned} \nabla U(\mathbf{x}_S) &= \mathbf{q}(\mathbf{x}_S) - \lambda \\ \lambda &\geq \mathbf{0} \\ \lambda^T \mathbf{x}_S &= 0 \end{aligned} \quad (12)$$

where the i 's element of $\mathbf{q}(\mathbf{x})$ is $x_i \cdot l'_i(x_i) + l_i(x_i)$.

If the i th element of \mathbf{x}_S , $x_{S,i} > 0$, then $\lambda_i = 0$. If $x_{S,i} = 0$, then the i th element of $\mathbf{q}(\mathbf{x}_S)$ is 0. Also, we clearly have $\frac{\partial U(\mathbf{x})}{\partial x_i} \geq 0, \forall \mathbf{x} \geq \mathbf{0}$, so λ_i must be 0. As a result, (12) can be simplified as

$$\nabla U(\mathbf{x}_S) = \mathbf{q}(\mathbf{x}_S) \quad (13)$$

To maximize the revenue (MR), on the other hand, the service provider solves

$$\begin{aligned} \max_{\mathbf{p}, \mathbf{x}} \quad & \mathbf{p}^T \mathbf{x} \\ \text{st} \quad & \mathbf{x} \text{ is an NE given } \mathbf{p} \end{aligned} \quad (14)$$

We claim that $\mathbf{p}^T \mathbf{x} = \mathbf{x}^T [\nabla U(\mathbf{x}) - \mathbf{l}(\mathbf{x})]$, where the i 'th element of $\mathbf{l}(\mathbf{x})$ is $l_i(x_i)$. Recall that at NE, \mathbf{x} solves (10), and that we have dropped the constraint $\sum_{i=1}^N d^i \leq D$ by assuming that the population is infinite. If $x_i > 0$, then $\frac{\partial U(\mathbf{x})}{\partial x_i} = p_i + l_i(x_i)$. So $x_i p_i = x_i [\frac{\partial U(\mathbf{x})}{\partial x_i} - l_i(x_i)]$. If $x_i = 0$ (i.e., it lies at the boundary of the feasible set), then $\frac{\partial U(\mathbf{x})}{\partial x_i} \leq p_i + l_i(x_i)$. But even in this case, $x_i [\frac{\partial U(\mathbf{x})}{\partial x_i} - l_i(x_i)] = 0 = x_i p_i$. Therefore, $\mathbf{p}^T \mathbf{x} = \mathbf{x}^T [\nabla U(\mathbf{x}) - \mathbf{l}(\mathbf{x})]$.

Denote the revenue-maximization solution as $\mathbf{x}_M \geq \mathbf{0}$. Then, the optimality condition is that for some $\lambda \geq \mathbf{0}$,

$$\begin{aligned} \nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) \cdot \mathbf{x}_M &= \mathbf{q}(\mathbf{x}_M) - \lambda \\ \lambda &\geq \mathbf{0} \\ \lambda^T \mathbf{x}_M &= 0 \end{aligned} \quad (15)$$

And the price (vector) charged by the SP is

$$\mathbf{p}_M = \nabla U(\mathbf{x}_M) - \mathbf{l}(\mathbf{x}_M)$$

Proposition 2: Define the price of anarchy $\rho = V_{MR}/V_{SO} = V(\mathbf{x}_M)/V(\mathbf{x}_S)$, which is the ratio between the social welfare at the revenue-maximizing point and at the social optimum point. Then we have a lower bound of ρ : $\rho \geq 2/3$, under either of the following sets of conditions:

- (a) Every element of $\nabla U(\mathbf{x})$ is concave in \mathbf{x} , and $\mathbf{x}_S \geq \mathbf{x}_M$ element-wise;
- (b) $(\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x})$ is concave in \mathbf{x} (at least for \mathbf{x} along the line from \mathbf{x}_M to \mathbf{x}_S). Note that condition (a) implies (b) and therefore is stronger.

Remark: The two concavity conditions are somewhat abstract. But they basically means that if the low-utility users are not dominant in the population, the efficiency loss due to the SP's selfish behavior is not too bad. (Similar assumptions are made in [7], [8] without modeling users' time preference.) For an opposite example, consider the curve in Fig. 1, which is actually $\nabla U(x) = U'(x)$ in the one-dimension case. It has a ‘‘long tail’’ instead of being concave, i.e., there are many low-utility users. Since the SP needs to charge a relatively high price to maximize its revenue, most users (which have low utilities) can not use the network, leading to a low value of social welfare and a ρ value close to 0. With other assumptions about the tail (not necessarily the concavity assumption made here), one may be able to find other bounds of POA.

The proof is in the Appendix. To give some intuition, Fig. 2 shows the simple case when there is only one time slot. The dashed curve is the marginal utility $U'(x)$, which is decreasing, and is also concave by the assumption in Proposition 2. The dotted curve is the marginal cost (delay) $q(x)$. Social Optimum is reached at point H , where $U'(x_S) = q(x_S)$ (by equation (13)). Assume that D is the point that gives the maximal

where the equality holds with $N = 1$. Indeed, ρ here is above $\frac{2}{3}$.

VI. CONCLUSION AND EXTENSIONS

In this paper, we have studied a new model for time-dependent pricing in the Internet. By likening different access times to different virtual links, the resulting model is similar to selfish routing and is amenable to analysis. However, different from most existing works on selfish routing, each user here has preference to different time slots. Under this model, we have shown that without pricing, the price of anarchy can be arbitrarily bad. From the SP's perspective, if he has perfect information about the user preferences, then differentiated pricing over users and times can result in maximal revenue and also social optimum (but all the benefits go to the SP). On the other hand, if the SP does not have user information except the amount of traffic at different time slots, it can only charge different prices over times. Then the resulting equilibrium can be far from social optimal. In general, the difference is unbounded; but under some conditions on the distribution of user preferences (specifically, if the low-utility users are not too many), the social welfare is at least 2/3 of the optimal value.

In this paper, we have assume that the network is composed of a single physical link (although time slots are modeled as "virtual links"). However, it can be extended to a network (owned by one SP) with multiple links inter-connected to each other, where the amount of traffic at one link is some linear combination of traffic from multiple routes, each of which traverses a number of links [10]. Assume that the users can choose both the time and physical routes to send/receive their traffic, then this is a selfish routing problem in two dimensions: times and physical routes. It is not difficult to see that the result in section III still holds: if each link charges the marginal cost according to equation (7), social optimum can be achieved. Also, the conclusion in section V-A (SP has full information) holds, and those in section V-B (SP has partial information) holds qualitatively, but possibly not quantitatively.

We have also assumed that each user has a fixed amount of traffic to send, and can only choose one time slot for simplicity. The model can be extended to the case where user i can use multiple time slots, and with different types of service (for example, voice or data). To do this, define a concave utility function, $u_i(\mathbf{s}_i)$, where $\mathbf{s}_i = (s_i^1, s_i^2, \dots, s_i^N) \in \mathcal{R}^N$ is the vector of the traffic volumes he sends/receives in each time slot. $u_i(\mathbf{s}_i)$ is not symmetric in the sense that given the vector \mathbf{s}_i , the marginal utility at slot j , $\frac{\partial u_i(\mathbf{s}_i)}{\partial s_j^i}$ is usually different for different j 's. This asymmetry reflects his time-preference. Also, $u_i(\mathbf{s}_i)$ takes different forms for different types of service. In this scenario, a social planner would still charge the marginal price to achieve social optimum. But when a monopolist SP tries to maximize its revenue with linear pricing, we expect that social optimum could not be achieved even if the SP knows the exact utility functions of the users, since the the revenue (price multiplied by volumn) is usually

not the same as (or "not aligned to") the user utility. A more detailed study of this scenario is left for future research.

This paper have focused on the analytical model. It is valuable to relate this model to empirical data on network usage or bandwidth trading. For example, the peak hours of network usage usually follows a certain pattern each day, and the revenue-maximizing SP is able to learn his optimal time-of-the-day prices over a period of time. So obtaining and comparing the data to our model is particularly interesting.

APPENDIX

Proof of Proposition 2:

Since condition (a) implies (b), we only prove the argument with condition (b). Define $L(\mathbf{x}) := \sum_{i=1}^N x_i \cdot l_i(x_i)$, then $V(\mathbf{x}) = U(\mathbf{x}) - L(\mathbf{x})$.

It's easy to see that $L(\mathbf{x})$ is convex, therefore $V(\mathbf{x})$ is concave. So,

$$V(\mathbf{x}_M) - V(\mathbf{0}) = V(\mathbf{x}_M) \geq \mathbf{x}_M^T \nabla V(\mathbf{x}_M)$$

We claim that

$$\begin{aligned} \mathbf{x}_M^T \nabla V(\mathbf{x}_M) &= \mathbf{x}_M^T [\nabla U(\mathbf{x}_M) - \mathbf{q}(\mathbf{x}_M)] \\ &= \mathbf{x}_M^T [-\nabla^2 U(\mathbf{x}_M) \cdot \mathbf{x}_M] \end{aligned} \quad (17)$$

For convenience, denote $\mathbf{s}(\mathbf{x}_M) := -\nabla^2 U(\mathbf{x}_M) \cdot \mathbf{x}_M$. First by definition, $\mathbf{x}_M^T \nabla V(\mathbf{x}_M) = \mathbf{x}_M^T [\nabla U(\mathbf{x}_M) - \mathbf{q}(\mathbf{x}_M)]$. If the i 'th element of \mathbf{x}_M , $x_{M,i} > 0$, then in (15), $\lambda_i = 0$. So $\frac{\partial U(\mathbf{x}_M)}{\partial x_i} - (\mathbf{q}(\mathbf{x}_M))_i = (\mathbf{s}(\mathbf{x}_M))_i$, making $x_{M,i} \cdot [\frac{\partial U(\mathbf{x}_M)}{\partial x_i} - (\mathbf{q}(\mathbf{x}_M))_i] = x_{M,i} \cdot (\mathbf{s}(\mathbf{x}_M))_i$. If $x_{M,i} = 0$, then $x_{M,i} \cdot [\frac{\partial U(\mathbf{x}_M)}{\partial x_i} - (\mathbf{q}(\mathbf{x}_M))_i] = x_{M,i} \cdot (\mathbf{s}(\mathbf{x}_M))_i = 0$. This proves the above claim.

So,

$$V(\mathbf{x}_M) \geq -\mathbf{x}_M^T \nabla^2 U(\mathbf{x}_M) \cdot \mathbf{x}_M$$

From the concavity assumption in condition (b), we have for any \mathbf{x} on the line segment between \mathbf{x}_M and \mathbf{x}_S ,

$$\begin{aligned} &(\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}) \\ &\leq (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) \\ &\quad + (\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M) (\mathbf{x} - \mathbf{x}_M) \end{aligned} \quad (18)$$

Now consider two integrals

$$\int_{\mathbf{x}_M}^{\mathbf{x}_S} [\nabla U(\mathbf{x})]^T d\mathbf{x} = U(\mathbf{x}_S) - U(\mathbf{x}_M)$$

and

$$\begin{aligned} &\int_{\mathbf{x}_M}^{\mathbf{x}_S} [\nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) (\mathbf{x} - \mathbf{x}_M)]^T d\mathbf{x} \\ &= (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) \\ &\quad + \frac{1}{2} (\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M) (\mathbf{x}_S - \mathbf{x}_M) \end{aligned}$$

If we perform the integration along the line from \mathbf{x}_M to \mathbf{x}_S , then $d\mathbf{x}$ is proportional to $\mathbf{x}_S - \mathbf{x}_M$. From (18), we have $[\nabla U(\mathbf{x})]^T d\mathbf{x} \leq [\nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) (\mathbf{x} - \mathbf{x}_M)]^T d\mathbf{x}$. So

$$\begin{aligned} &U(\mathbf{x}_S) - U(\mathbf{x}_M) \\ &\leq (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) \\ &\quad + \frac{1}{2} (\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M) (\mathbf{x}_S - \mathbf{x}_M) \end{aligned}$$

It follows that

$$\begin{aligned}
& V(\mathbf{x}_S) - V(\mathbf{x}_M) \\
&= U(\mathbf{x}_S) - U(\mathbf{x}_M) - [L(\mathbf{x}_S) - L(\mathbf{x}_M)] \\
&\leq (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) + \frac{1}{2}(\mathbf{x}_S - \mathbf{x}_M)^T \cdot \\
&\quad \nabla^2 U(\mathbf{x}_M)(\mathbf{x}_S - \mathbf{x}_M) - (\mathbf{x}_S - \mathbf{x}_M)^T \mathbf{q}(\mathbf{x}_M)
\end{aligned}$$

Using (17) and (15), we have

$$\begin{aligned}
& V(\mathbf{x}_S) - V(\mathbf{x}_M) \\
&\leq (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) + \\
&\quad \frac{1}{2}(\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M)(\mathbf{x}_S - \mathbf{x}_M) \\
&\quad - \mathbf{x}_S^T \mathbf{q}(\mathbf{x}_M) + \mathbf{x}_M^T [\nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M] \\
&\leq (\mathbf{x}_S - \mathbf{x}_M)^T \nabla U(\mathbf{x}_M) + \\
&\quad \frac{1}{2}(\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M)(\mathbf{x}_S - \mathbf{x}_M) \\
&\quad - \mathbf{x}_S^T [\nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M] + \\
&\quad \mathbf{x}_M^T [\nabla U(\mathbf{x}_M) + \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M] \\
&= -(\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M + \\
&\quad \frac{1}{2}(\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M)(\mathbf{x}_S - \mathbf{x}_M)
\end{aligned}$$

Then, consider

$$\begin{aligned}
& V(\mathbf{x}_M) - 2[V(\mathbf{x}_S) - V(\mathbf{x}_M)] \\
&\geq -\mathbf{x}_M^T \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M + \\
&\quad (\mathbf{x}_S - \mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M)(3\mathbf{x}_M - \mathbf{x}_S) \\
&= -4\mathbf{x}_M^T \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M + 4\mathbf{x}_S^T \nabla^2 U(\mathbf{x}_M) \mathbf{x}_M \\
&\quad - \mathbf{x}_S^T \nabla^2 U(\mathbf{x}_M) \mathbf{x}_S \\
&= -(\mathbf{x}_S - 2\mathbf{x}_M)^T \nabla^2 U(\mathbf{x}_M)(\mathbf{x}_S - 2\mathbf{x}_M) \\
&\geq 0
\end{aligned}$$

where the last step follows from the fact that $U(\mathbf{x})$ is concave. Therefore,

$$\rho = \frac{V(\mathbf{x}_M)}{V(\mathbf{x}_S)} \geq \frac{2}{3}$$

REFERENCES

- [1] Tim Roughgarden, "The price of anarchy is independent of the network topology", Proceedings of the thirty-fourth annual ACM symposium on Theory of computing, Montreal, Quebec, Canada, Pages: 428 - 437, Year: 2002
- [2] Tim Roughgarden, "How bad is selfish routing", Journal of the ACM, 2002
- [3] J. G. Wardrop, "Some theoretical aspects of road traffic research," in Proceedings of the Institute of Civil Engineers, vol. 1, 1952, pp. 325-378.
- [4] E. Koutsoupias, C. H. Papadimitriou, "Worst-case equilibria," Annual Symposium on Theoretical Aspects of Computer Science, 1999.
- [5] Daron Acemoglu and Asuman Ozdaglar, "Flow Control, Routing, and Performance from Service Provider Viewpoint", LIDS report, MIT, 2003
- [6] Daron Acemoglu and Asuman Ozdaglar, "Competition and Efficiency in Congested Markets", Mathematics of Operations Research, 2007
- [7] Asuman Ozdaglar, "Price Competition with Elastic Traffic", LIDS report, MIT, 2006
- [8] John Musacchio and Shuang Wu, "The Price of Anarchy in a Network Pricing Game", Allerton Conference 2007
- [9] Ioannis Ch. Paschalidis, John N. Tsitsiklis, "Congestion-Dependent Pricing of Network Service", Networking, IEEE/ACM Transactions on, Volume 8, Issue 2, Apr 2000, Page(s):171 - 184

- [10] F. P. Kelly, A.K. Maulloo and D.K.H. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," Journal of the Operational Research Society 49 (1998), 237-252.
- [11] S. Boyd and L. Vandenberghe, "Convex Optimization", Cambridge University Press, 2004.