The Case for Non-cooperative Multihoming of Users to Access Points in IEEE 802.11 WLANs

Srinivas Shakkottai
Dept. of Electrical and Computer Engineering, and
Coordinated Science Laboratory
University of Illinois at Urbana-Champaign, USA
Email: sshakkot@uiuc.edu

Eitan Altman
INRIA
Sophia Antipolis Cedex
France
Email: altman@sophia.inria.fr

Anurag Kumar
Dept. of Electrical Communication Engineering
Indian Institute of Science
Bangalore, India
Email: anurag@ece.iisc.ernet.in

Abstract— **In many cases, a mobile user has the option of connecting to one of several IEEE 802.11 access points (APs), each using an independent channel. User throughput in each AP is determined by the number of other users as well as the frame size and physical rate being used. We consider the scenario where users could *multihome*, i.e., split their traffic amongst all the available APs, based on the throughput they obtain and the price charged. Thus, they are involved in a non-cooperative game with each other. We convert the problem into a fluid model and show that under a pricing scheme, which we call the cost price mechanism, the total system throughput is maximized, i.e., the system suffers no loss of efficiency due to selfish dynamics. We also study the case where the Internet Service Provider (ISP) could charge prices greater than that of the cost price mechanism. We show that even in this case multihoming outperforms unihoming, both in terms of throughput as well as profit to the ISP.**

I. Introduction

The IEEE 802.11 protocol is currently the standard for wireless LANs (WLANs), with no fundamental difference between the different flavors. It has been deployed ubiquitously in airports, coffee shops and homes. Very often there is a choice of access points (APs) to which a mobile user could connect to. Users scan the wireless channel in order to find the AP which shows the highest signal strength and associate to it. They then transmit at different rates (often called the PHY rate) based on the signal strength indicated. The algorithm that selects the PHY rate chooses a higher rate if the signal strength is good and progressively cuts down the rate as signal strength decays. It achieves such rate adaptation by keeping the transmit power almost constant, while changing the constellation used. Thus, it would use BPSK for a bad channel, QPSK for a better one and so on. But this also means that for a frame transmission of the same size, some users occupy the channel longer than others. It has also been observed [1] that all the connections in a single cell receive the same throughput, leading to inefficient use of the channel. In such a scenario, the question arises as to whether it might be better for a user to split his or her traffic among the visible APs.

Suppose we have a geographical region divided into cells as shown in Figure 1. Each cell would have an access point. Transmissions in each cell would be independent of other cells by using separate channels. For example 802.11 b and g have three independent channels and we may use them to tessellate a region into independent cells. Another possible scenario is when the same region has multiple independent access points, perhaps provided by competing service providers. In either case, users might have the option of connecting to one of several access points based on where they are located. For instance, in Figure 1, users in region A might be able to associate to cells P, Q and S, whereas users in region B, might have no choice but to associate to cell S.

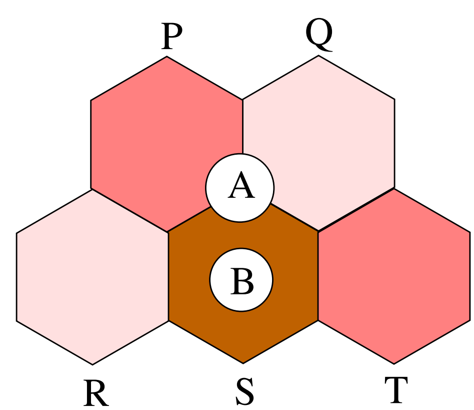

Fig. 1. Division of a geographical region into non-interfering cells using three independent channels, indicated by shading. Users could be in a position to connect to access points in one or more cells.

Users with just a single wireless network interface card might also be able to associate to all the APs available to them, which would provide diversity from the fact that different cells may be loaded differently. They could then probabilistically divide their traffic among the different APs in order to maximize their individual throughput. Virtualizing a wireless card in this manner has actually been implemented [2]. Traffic splitting in the Internet among different Internet Service Providers (ISPs) is called *multihoming* [3] and we follow the same terminology for the WLAN case. We call the case where users can associate to only a single AP as *unihoming*. Of course, in our case all the APs might be owned by the same ISP. We assume that users are aware of the throughput that they would obtain if they joined one of the APs (they would have to run an estimation tool using a test sequence of packets or the AP could provide the current

1-4244-0222-0/06/$20.00 (c)2006 IEEE

system state). This would tell them the potential benefit if they sent traffic to that AP. The AP itself might charge a price for sending packets through it. So the payoff that the user obtains would be the difference of the two. We also assume that users do not have the freedom to choose frame sizes or PHY rates as they wish – they are decided by the operating system.

Users are selfish and would like to maximize their payoffs. Thus, they compete with each other in a non-cooperative game.

Related Work

There has recently been much interest in understanding the behavior of wireless LANs. They make use of the distributed coordination function (DCF) with an RTS-CTS handshake and hence cannot directly be modeled in the same manner as traditional Ethernet systems. One intriguing question has been that of why users using different PHY rates all obtain the same throughput. This question was studied using simulation and experiments in [1]. In [4] the system was studied as a two-player game, with each user trying to maximize their individual throughput and results were presented on the inefficiency of the system as compared to the cooperative optimum. Bianchi [5] used fixed point analysis in order to provide an analytical framework for 802.11 WLANs. The results were extended in [6], to provide expressions for the throughput of users with disparate frame sizes and PHY rates. Our work relies heavily on the expressions obtained in the above. The analytical work has been further extended in [7] and a simulation based verification provided.

Another area that has received attention is that of how users should associate to APs in a WLAN. In [8] a study is made on fairness issues and how the load should be balanced using fractional association in a cooperative scenario. Usually, users have no particular incentive to cooperate with each other and would be interested in maximizing their individual payoffs. In [9], the case of non-cooperative users who decide on the optimal frame size and PHY rate to be used in order to maximize their individual throughputs is studied. The users are all assumed to be in a single cell and compete for throughput within that cell. Another paper on non-cooperative association is [10], which provides a simulation study of the benefit of associating to the AP that would provide the best estimated link rate. Some results on cooperative association of users to different APs are provided in [11].

Multihoming is a recent idea that has been proposed to make use of path diversity in the Internet. The idea is that since different ISPs use different policy based routing mechanisms, it is very possible that a user would get a higher bandwidth by subscribing to multiple ISPs simultaneously and splitting traffic among them. Another concept which achieves the same at a finer resolution is that of source routing, wherein the user chooses the routes by himself, rather than choosing ISPs. A comparative study of overlay source routing and multihoming is carried out in [3]. One question which crops up when multihoming is allowed is that of how users ought to split up their traffic among the different ISPs. A dynamic programming algorithm based on how much different ISPs charge is studied in [12], where it is assumed that the ISPs have sufficient capacity to handle the traffic at an acceptable throughput for the users. Analytical work on the stability of a system using multiple routes is present in [13], [14]. The first studies a multi-path TCP version, which would split traffic among the different routes, as a feedback system with delays and finds the required gain for stability. The second studies a general class of decentralized algorithms that would optimally split traffic.

Recently there has been an attempt to extend the multihoming idea to wireless LANs. The idea of virtualizing a wireless card has been studied in [2]. The authors use the sleep feature of 802.11 cards in order to switch between APs. The idea is to make one AP believe that the wireless card is asleep (which would cause the AP to buffer packets), while actually sending traffic to another AP. They also propose empirical methods by which the ratio of traffic associated to each AP could be chosen.

Selfish routing and multihoming bring issues of system efficiency with them. A completely centralized scheme could, in theory, optimize the system throughput. However, this kind of control is usually not feasible. By providing a choice for the users, one increases the anarchy of the system. Then the question immediately arises as to whether Nash/Wardrop equilibria exist, and how much efficiency loss occurs due to this anarchy. Analytical studies of this sort are available in [15]–[19] and provide bounds on the worst case efficiency. In [20] measurement traces on the Internet are used to study the effects of selfish routing.

In many studies of traffic using selfish routing, one would like to think of users, not as integral values, but as real numbers. The reason for this is usually because the number of users is large (for example, in modeling highways or backbone Internet fibers). A concept that has been applied successfully to obtain quantitative results is that of the Wardrop equilibrium [21]. A comprehensive description of traffic models using the concept of infinitesimal users is present in [22].

Internet pricing is a topic of considerable interest today. Clearly, any scheme however efficient cannot be implemented unless it is worthwhile for the ISPs to do so. One example of differentiated pricing to provide different perceived QoS is Paris Metro Pricing (PMP) [23], which is also studied in [24]. Some examples of literature that deals with pricing strategies and competition on the Internet are [25]–[27].

Our study builds upon and extends the above work. We study multihoming in a relatively new arena – that of WLANs – with its own array of attendant issues. Particular to 802.11 is the fact that the throughput of the system is not fixed, but depends on the distribution of user types. Another interesting fact is that (assuming that frame sizes are fixed), the throughput of all the users, regardless of their PHY rates is the same. Our contributions are detailed in the following subsection.

Main Results

We consider the expressions for the many users regime obtained in [6], and use it to construct a fluid model of user masses which can multihome to different APs. We allow users to use mixed strategies, i.e., they choose alternatives probabilistically. The deterministic equivalent of this situation is that user masses would split among the alternatives, with the mass being proportional to the probability of choosing that particular option. Thus, the ratio in which the masses are divided amongst the different APs gives the probabilities of associating with them. For example, if 3 units of a class of users are associated to one AP and 1 unit to another AP, it would mean that the strategy that the class of users play is $[\frac{3}{4}\frac{1}{4}]$. This provides a framework in which selfish movement of user masses can be studied deterministically. We thus transform the problem to that of a *population game*, which is designed for the study of such non-cooperative systems.

In the WLAN scenario, intuitively it seems clear that since different users send at different PHY rates, their "occupancy" of the channel is different. We formalize the idea of occupancy, and propose a pricing mechanism in which users are charged based on their channel occupancy. We call this "cost price charging". The difference of the throughput and the price charged gives the payoff to the users. We study the game under the assumption that at a given time users would try to take that action which is most profitable. Descriptions of two such dynamics exist in game theoretical literature – replicator dynamics [28] and Brown-von Neumann-Nash dynamics [29]. Using the theory of Lyapunov functions, Sandholm [17] provides results on the conditions required for a general class of dynamics to be asymptotically stable. He also shows that such equilibria could be Wardrop equilibria. Following his techniques, we first verify that both types of dynamics satisfy the required conditions and find a suitable Lyapunov function in order to confirm stability under the cost price mechanism. We show that the payoffs at equilibrium in each cell in use by a particular class of users are all equal. The solution so obtained would be a Wardrop equilibrium [21], consistent with Sandholm's results

We next turn to characterizing the nature of the equilibrium. We would like to know how much efficiency loss is suffered due to decentralized, selfish multihoming. This would tell us the price of anarchy for the 802.11 WLAN system. Again, Sandholm [17] has a result that states there is no loss of efficiency under strong symmetry conditions. Although these conditions are not satisfied in our scenario, we show that there is no loss of efficiency due to selfish multihoming, i.e., *anarchy is obtained at no cost*. This is interesting since it essentially says that multihoming in WLANs is ideal for decentralized control. Charging users the cost price of their occupancy causes them to split their masses optimally.

Finally, we deal with the economics of multihoming – whether or not it makes sense economically for an ISP to permit multihoming in its APs. We show that when an ISP charges differentiated prices above the cost price charge in the different APs, multihoming achieves at least the same profit as unihoming. So the ISP suffers no loss by allowing its customers to multihome. We further show that even in the case of differentiated pricing, the throughput of the system as a whole is at least that of unihoming, thus building a strong case for multihomed IEEE 802.11 wireless LANs.

Organization of the Paper

The paper is organized as follows. In Section II we discuss the game theoretic concepts used. We then discuss the required background on 802.11 WLANs in Section III. The section presents the expressions derived in [6] that are relevant to this work. In Section IV, we specify the model of the WLAN with multiple classes of users and present its fluid equivalent. We then proceed in Section V, to study the dynamics of the system in a non-cooperative scenario. The idea here is to show that the system is stable using Lyapunov techniques. We next study the efficiency of such an equilibrium in Section VI and show that the Wardrop equilibrium is efficient. We study the economic impact of multihoming in Section VII. We show that allowing users to multihome does not hurt profits and that even under differentiated pricing, multihoming outperforms unihoming in terms of throughput. We also briefly discuss price selection and conclude with pointers to extensions in Section VIII.

II. Basic Ideas on Population Games

We first introduce the game theoretic concepts that are used in this paper. A good reference on game theory is [30], and much of the discussion below may be found in [17]. A *population game* F, with Q non-atomic classes of players is defined by a mass and a strategy set for each class and a payoff function for each strategy. By a non-atomic population, we mean that the contribution of each member of the population is infinitesimal. We denote the set of classes by $\mathcal{Q} = \{1, ..., Q\}$, where $Q \geq 1$. The class $q \in \mathcal{Q}$ has mass $\hat{d}_q$. The set of strategies for class q is denoted $\mathcal{S}_q = \{1, ..., S_q\}$. These strategies can be thought of as the actions that members of q could possibly take. A particular strategy distribution is the way the class q partitions itself into the different actions available, i.e., a strategy distribution for q is vector of the form $y_q = \{y_q^1, y_q^2, ...y_q^{S_q}\}$, where $\sum_{i=1}^{S_q} y_q^i = \hat{d}_q$. The set of strategy distributions of a class $q \in \mathcal{Q}$, is denoted by $Y_q = \{y_q \in \mathbb{R}_+^{S_q} : \sum_{i=1}^{S_q} y_q^i = \hat{d}_q\}$. We denote the vector of strategy distributions being used by the entire population by $\mathbf{y} = \{y_1, y_2, ..., y_Q\}$, where $y_i \in Y_i$. The vector $\mathbf{y}$ can be thought of as the state of the system. Let the space of all strategy distributions be $\mathcal{Y}$.

The marginal payoff function (per unit mass) obtained from strategy $i \in \mathcal{S}_q$ by users of class q, when the state of the system is $\mathbf{y}$ is denoted by $F_q^i(\mathbf{y}) \in \mathbb{R}$ and is assumed to be continuous and differentiable. Note that the payoffs to a strategy in class q can depend on the strategy distribution within class q itself. The total payoff to users of class q is then given by $\sum_{i=1}^{S_q} F_q^i(\mathbf{y}) y_q^i$. Players may be cooperative or non-cooperative in behavior.

A commonly used concept in non-cooperative games is that of the Nash equilibrium. A particular state $\mathbf{y}$ is a Nash equilibrium if no unilateral deviation can allow the deviator to strictly gain. Whereas the Nash equilibrium is the right concept for the case of atomic players, in the context of infinitesimal

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the Proceedings IEEE Infocom.

players, a more appropriate idea is the Wardrop equilibrium [21]. Consider any strategy distribution $y_q = [y_q^1, ..., y_q^{S_q}]$. There would be some elements which are non-zero and others which are zero. We call the strategies corresponding to the non-zero elements as the *strategies used by class* q.

Definition 1 A state $\hat{\mathbf{y}}$ is a Wardrop equilibrium if for any class $q \in \mathcal{Q}$, all strategies being used by the members of q yield the same marginal payoff to each member of q, whereas the marginal payoff that would be obtained by members of q is lower for all strategies not used by class q.

Let $\hat{\mathcal{S}}_q \subset \mathcal{S}_q$ be the set of all strategies used by class q in a strategy distribution $\hat{\mathbf{y}}$. A Wardrop equilibrium $\hat{\mathbf{y}}$ is then characterized by the following relation:

$$F_q^s(\hat{\mathbf{y}}) \geq F_q^{s'}(\hat{\mathbf{y}}) \ \ \forall s \in \hat{\mathcal{S}}_q \text{ and } s' \in \mathcal{S}_q$$

■

One question which is important in identifying such Wardrop equilibria is that of the population dynamics that would lead to Wardrop equilibria. If each class q follows some dynamics, then would the stationary points be Wardrop equilibria? We present a result from [17], which is useful in this regard. We first need the following definition:

Definition 2 The dynamics $\dot{\mathbf{y}} = \mathbf{V}(\mathbf{y})$ are said to be *positively correlated* (PC) if

$$\sum_{k=1}^{Q} \sum_{i=1}^{S_k} F_k^i(\mathbf{y}) V_k^i(\mathbf{y}) > 0 \text{ whenever } V(\mathbf{y}) \neq 0$$

■

Result 1 If $\mathbf{V}(\mathbf{y})$ satisfies PC, all Wardrop equilibria of F are the stationary points of $\dot{\mathbf{y}} = \mathbf{V}(\mathbf{y})$. ■

Potential games are a subclass of games that have a specific structure on the cost. This structure allows to obtain convergence to equilibrium for various dynamics. The theory behind them is very similar to the theory of Lyapunov functions in control systems. The idea is to identify a scalar function which is used to represent the potential of the system. Users would try to maximize their payoffs at each time instant, thus raising the potential of the system. Using such a function it may be possible to show that a system of players, each following his or her own selfish dynamics, actually converges to a Wardrop equilibrium.

Definition 3 We call F a potential game if $\exists$ a C^1 function $\mathcal{T} : \mathcal{Y} \to \mathbb{R}$ such that $\frac{\partial \mathcal{T}}{\partial y_q^i}(\mathbf{y}) = F_q^i(\mathbf{y})$ for all $\mathbf{y} \in \mathcal{Y},\ i \in \mathcal{S}_q$ and $q \in \mathcal{Q}$. ■

The definition says that the rate of change of potential with mass of a population is the payoff obtained per unit mass by that population at any state. We then immediately have that if F is a potential game and $\mathbf{V}(\mathbf{y})$ is PC, then the potential function $\mathcal{T}$ is a Lyapunov function for the system $\dot{\mathbf{y}} = \mathbf{V}(\mathbf{y})$. This means that all the stationary points of $\dot{\mathbf{y}} = \mathbf{V}(\mathbf{y})$ would be asymptotically stable. Thus, we have the following useful result:

Result 2 A potential game F, with dynamics $\mathbf{V}(\mathbf{y})$ that are PC, has asymptotically stable stationary points. ■

In accordance with Result 1, the system state would converge to either a Wardrop equilibrium or a boundary point of the set $\mathcal{Y}$.

System Dynamics

We introduce two expressions commonly used to model population dynamics and show below that both of them are positively correlated. We also show that a combination of the two is PC.

The first dynamics that we study is called *Replicator Dynamics* [28]. The rate of increase of $\dot{y}_q^s / y_q^s$ of the strategy s is a measure of its evolutionary success. Following the basic tenet of Darwinism, we may express this success as the difference in fitness $F_q^s(\mathbf{y})$ of the strategy s and the average fitness $\sum_{i=1}^{S_q} y_q^i F_q^i(\mathbf{y}) / \hat{d}_q$ of the class q. Then we obtain

$$\frac{\dot{y}_q^s}{y_q^s} = \text{fitness of } s \text{ - average fitness.}$$

Then the dynamics used to describe changes in the mass of class q playing strategy s is given by

$$\dot{y}_q^s = \mathbf{V}(\mathbf{y}) = y_q^s \left(F_q^s(\mathbf{y}) - \frac{1}{\hat{d}_q} \sum_{i=1}^{S_q} y_q^i F_q^i(\mathbf{y}) \right). \tag{1}$$

Note that the dynamics take place within the set $\sum_{j=1}^{S_q} y_q^j = \hat{d}_q \ \ \forall q \in \{1, 2, ..., Q\}$, i.e., the total mass of each class remains fixed. This fact may be seen immediately by summing (1) over all strategies, yielding $\sum_{s=1}^{S_q} \dot{y}_q^s = 0$. The above expression thus says that a population would increase the mass of a successful strategy and decrease the mass of a less successful one. It is called the replicator equation after the tenet "like begets like". We show below that replicator dynamics satisfy PC. We provide a proof as we could not find one in the literature reviewed.

Theorem 1: The system with replicator dynamics is positively correlated.

Proof:

$$\sum_{k=1}^{Q} \sum_{i=1}^{S_k} F_k^i(\mathbf{y}) V_k^i(\mathbf{y}) = \sum_{k=1}^{Q} \sum_{i=1}^{S_k} F_k^i(\mathbf{y}) \frac{\partial y_k^i}{\partial t}$$

From (1), we have

$$= \sum_{k=1}^{Q} \sum_{i=1}^{S_k} F_k^i(\mathbf{y}) y_k^i \left(F_k^i(\mathbf{y}) - \frac{1}{\hat{d}_k} \sum_{j=1}^{S_k} y_k^j F_k^j(\mathbf{y}) \right)$$

$$= \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} \frac{y_k^i}{\hat{d}_k} \left(F_k^i(\mathbf{y}) \right)^2 - \left(\sum_{i=1}^{S_k} \frac{y_k^i}{\hat{d}_k} F_k^i(\mathbf{y}) \right)^2 \right)$$

Now, since $\sum_{j=1}^{S_k} \frac{y_k^j}{\hat{d}_k} = 1$, by Jensen's inequality we have that the term in parentheses above is non-negative (with equality at $\mathbf{V}(\mathbf{y}) = 0$). Thus, the summation is also non-negative and the proof follows. ■

Another commonly used model is called *Brown-von Neumann-Nash (BNN)* dynamics [29], which is somewhat more complex. Let,

$$\gamma_q^s = \max\left\{F_q^s(\mathbf{y}) - \frac{1}{\hat{d}_q}\sum_{i=1}^{S_q} y_q^i F_q^i(\mathbf{y}), 0\right\} \quad (2)$$

denote the excess marginal payoff to strategy s relative to the average payoff in its class. Then BNN dynamics are described by

$$\dot{y}_q^s = \mathbf{V}(\mathbf{y}) = \hat{d}_q \gamma_q^s - y_q^s \sum_{j=1}^{S_q} \gamma_j^s, \quad (3)$$

where the dynamics take place within the set $\sum_{j=1}^{S_q} y_q^j = \hat{d}_q \quad \forall q \in \{1, 2, ..., Q\}$. An interpretation of the BNN dynamics is that during any short time interval, all players in a class are equally likely to switch strategies, and do so at a rate proportional to the sum of the excess payoffs in the class. Those who switch choose strategies with above average payoffs, choosing each with probability proportional to the strategy's excess payoff. The proof that BNN dynamics are PC is present in [17], but we repeat it here for completeness.

Theorem 2: The system with BNN dynamics is positively correlated.

Proof: Define $\bar{F}_q \triangleq \frac{1}{\hat{d}_q}\sum_{i=1}^{n_q} y_q^i F_q^i(\mathbf{y})$. Then we have

$$\begin{aligned}
\sum_{k=1}^{Q}\sum_{i=1}^{S_k} F_k^i(\mathbf{y}) V_k^i(\mathbf{y}) &= \sum_{k=1}^{Q}\sum_{i=1}^{S_k} F_k^i(\mathbf{y}) \frac{\partial y_k^i}{\partial t} \\
= \quad & \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} F_k^i(\mathbf{y})\gamma_k^i - \frac{1}{\hat{d}_k}\sum_{i=1}^{S_k} y_k^i F_k^i(\mathbf{y}) \sum_{j=1}^{S_k}\gamma_k^j\right) \\
= \quad & \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} F_k^i(\mathbf{y})\gamma_k^i - \bar{F}_k \sum_{j=1}^{S_k}\gamma_k^j\right) \\
= \quad & \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} F_k^i(\mathbf{y})\gamma_k^i - \sum_{j=1}^{S_k}\gamma_k^j \bar{F}_k\right) \\
= \quad & \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} \gamma_k^i \left(F_k^i(\mathbf{y}) - \bar{F}_k\right)\right) \\
= \quad & \sum_{k=1}^{Q} \hat{d}_k \left(\sum_{i=1}^{S_k} \left(\gamma_k^i\right)^2\right) \geq 0 \text{ (with equality at } \mathbf{V}(\mathbf{y}) = 0).
\end{aligned}$$

Hence the proof. ■

The reason for considering BNN dynamics is that unlike replicator dynamics, it has the property of non-complacency in that it allows extinct strategies to resurface, so that its stationary points are always Wardrop equilibria [17].

So far we considered a class to be population of users with the same set of available strategies. Suppose we expand the definition of "class" to also include the dynamics being followed. So a class now a population of users with the same set of strategies and following the same dynamics. Then it is straightforward to prove the following corollary.

Corollary The system is PC as long as each class follows either replicator or BNN dynamics.

Proof: The proof is simple. We have

$$\frac{\partial T(\mathbf{y}(t))}{\partial t} = \sum_{k=1}^{Q}\sum_{i=1}^{S_k} \frac{\partial T(\mathbf{y}(t))}{\partial y_k^i}\frac{\partial y_k^i}{\partial t}$$

Define

$$\zeta_k(\mathbf{y}) \triangleq \sum_{i=1}^{S_k} \frac{\partial T(\mathbf{y}(t))}{\partial y_k^i}\frac{\partial y_k^i}{\partial t}.$$

Then from the proofs of Theorem's 1 and 2, we immediately have $\zeta_k(\mathbf{y}) \geq 0$ for all $k \in \{1, ...Q\}$, from which the proof follows. ■

We have thus shown that under two standard models of selfish dynamics (or a combination thereof), the system is positively correlated.

III. Background on IEEE 802.11 WLANs

We provide the relevant background on expressions relating to the throughput of an IEEE 802.11 cell.

The single cell

We begin by recalling uplink throughput expressions for a cell containing a single AP obtained in [6], [9]. It holds when the nearest co-channel AP is farther away than the carrier-sense range (as we assumed in Figure 1). We use this throughput as a measure of the payoff derived from associating to a particular AP. The expressions are for the MAC layer. Let there be n active users in a *single cell* IEEE 802.11 WLAN contending to transmit data. Each user uses the Distributed Coordination Function (DCF) protocol with an RTS/CTS frame exchange before any data-ack frame exchange and has an equal probability of the channel being allocated to it. It is assumed that every user has infinitely many packets backlogged in its transmission buffer. In other words, the transmission buffer of each user is *saturated* in the sense that there are always packets to transmit when a user gets a chance to do so. It is also assumed that all the users use the same *back-off* parameters. Let β denote the long run average attempt rate per user per slot ($0 \leq \beta \leq 1$) in *back-off time* [1] (Conditions for the existence of a unique such β are given in [7].)

Call the cell s. Let the MAC frame size of user i be L_i bits and let the PHY rate used by this user be denoted by R_i^s bits per slot. Let T_o be defined as the transmission overhead in slots related to a frame transmission, which comprises of the SIFS/DIFS, etc and let T_c be defined as the fixed overhead for an RTS collision in slots. Then it follows from [6] that the throughput of user i is given by

$$\theta(i,n) = \frac{\beta e^{-n\beta} L_i}{1 + n\beta e^{-n\beta}\left(T_o - T_c + \frac{1}{n}\sum_{i=1}^{n}\frac{L_i}{R_i^s}\right) + \left(1 - e^{-n\beta}\right)T_c},$$

[1] If we plot transmission attempts as a function of "real" time, and then *cut out* from the plot the channel activity periods (during which all users freeze their back-off), then the new horizontal axis is called the "back-off time", see Section II.A of [6].

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the Proceedings IEEE Infocom.

where $\beta = \beta(n)$ (i.e. β is a function of n) is obtained as the solution of a fixed point equation that does not depend on L_i's or R_i^s's. As is the case in IEEE 802.11, for all users that use an RTS/CTS frame exchange before the data-ack frame transmission, we assume throughout our discussion that

$$T_o \geq T_c$$

To find the limit as $n \to \infty$, we identify here the asymptotic aggregate throughput as $n \to \infty$ An appealing feature of the asymptotic case is that we have an *explicit* expression for $\beta(n)$.

Asymptotic throughput

Let p be the exponential back-off multiplier, i.e. if b_z is the mean back-off duration (in slots) at the zth attempt for a frame then $b_z = p^z b_0$. According to the IEEE 802.11 specifications $p = 2$. Each user uses one of the Q distinct available values of the parameters (L_i, R_i^s) with $(L_i, R_i^s) \in \{(L_1, R_1^s), ..., (L_Q, R_Q^s)\}$. We derive the corresponding asymptotic throughput. Assume that there are m_q users using parameters (L_q, R_q^s). Denote by $\alpha_q(n) = m_q/n$ the fraction of the users using (L_q, R_q^s) among all users in the cell. Then the throughput of all users using (L_q, R_q^s) is given by

$$\tau(\alpha_q(n)) = \frac{m_q \beta e^{-n\beta} L_q}{1 + n\beta e^{-n\beta}\left(T_o - T_c + \sum_{i=1}^{Q} \frac{\alpha_i(n) L_i}{R_i^s}\right) + (1 - e^{-n\beta})\, T_c} \quad (4)$$

It is assumed that $\alpha_q(n)$ converges to a limit α_q. Note that the attempt rate $\beta = \beta(n)$ and the collision probability are not functions of L_i nor of R_i^s. As in [6], taking the limit as $z \to \infty$ and $n \to \infty$, it can be observed that

$$\lim_{n\to\infty} n\beta(n) \uparrow \ln\left(\frac{p}{p-1}\right), \quad (5)$$

where $\beta(n)$ is obtained as the solution of a fixed point equation corresponding to n users (see Theorem VII.2 in [6]). Combining (4) and (5) we get as $n \to \infty$ the following expression for the aggregate throughput of all users using (L_q, R_q^s):

$$\tau(\alpha_q) = \frac{\alpha_q L_q}{\kappa + \sum_{j=1}^{Q} \frac{\alpha_j L_j}{R_j^s}}, \quad (6)$$

where

$$\kappa = \frac{p + T_c}{(p-1)\ln\left(\frac{p}{p-1}\right)} + T_o - T_c \quad (7)$$

IV. System Model

Let there be S independent APs, which use different channels and so do not interfere with each other. We define a class q of users as the set of all users that have access to the same APs and common values of $[L_q, R_q^1, R_q^2, ..., R_q^{S_q}]$. Here S_q is the number of APs available to users of class q, L_q is the frame size, and R_q^i is the PHY rate that a user of class q would have if it connected to the ith AP. In 802.11, the PHY rate remains constant at the order of seconds. We assume that at the timescale being considered, users do not change their PHY rates. The class is used to model the fact that users in the same geographical location would face a similar set of circumstances. For instance, in Figure 1 users in region A would belong to a different class than users in region B. Let Q be the number of such classes. Thus, all users in a class q would have an identical set of options open to them. Let the users be capable of multihoming. Then their strategies consist of probability vectors of associating to each AP available to them.

Fluid Model

We wish to study the effects of the movement of masses of individuals of each class on their individual payoffs in a deterministic fashion. In order to do this we would like to consider users as infinitesimally divisible, i.e., consider a fluid model. Since all the expressions are in terms of integral quantities, we scale the system by letting $n \to \infty$, i.e., we consider the case where the number of users gets large. When scaling the number of users in this fashion, we must preserve the relative presence of each class of users in the whole. We then have a model, wherein different classes of users can distribute their masses amongst the different available APs. As before, a particular strategy distribution is the way the population partitions itself among the different APs available. As mentioned in the introduction, the ratio in which the masses are divided amongst the different APs is proportional to the probabilities of associating with them. Note that all the payoffs would be in the expected sense, i.e., users would actually be considering the expected payoff of assigning particular probabilities to the different APs. Thus, we convert a probability model with integral players into a deterministic fluid model, for which we study Wardrop equilibria. Although the Wardrop equilibrium is defined for infinitesimal players, it is known to be a good approximation for the case of atomic players under mild conditions, provided their number is sufficiently large [31].

Let there be d_q users of class q. Of these, assume a fraction x_q^s is connected to AP s. The total number of users connected to AP s is then given by

$$n^s = \sum_{q=1}^{Q} d_q x_q^s.$$

We define

$$\alpha_q^s \triangleq \frac{d_q x_q^s}{n^s} = \frac{d_q x_q^s}{\sum_{i=1}^{Q} d_i x_i^s},$$

which is understood to be zero if the denominator is zero. We wish to take the limit as n^s becomes large simultaneously for all s as a common parameter n goes to infinity, thus keeping the fractions fixed as we scale n. Hence, we consider the following scaling:

$$d_q = n\hat{d}_q.$$

n can be interpreted as the sum of all demand, i.e., $n = \sum_{q=1}^{Q} d_q$. As $n \to \infty$, we get from (6) that the throughput received by the total mass of users of class q connected to AP s is

$$\tau_q^s(\mathbf{x}^s) = \frac{\frac{L_q \hat{d}_q x_q^s}{\sum_{i=1}^{Q} \hat{d}_i x_i^s}}{\kappa + \sum_{j=1}^{Q} \frac{\left(\frac{\hat{d}_j x_j^s L_j}{R_j^q}\right)}{\sum_{i=1}^{Q} \hat{d}_i x_i}} \tag{8}$$

The term $\hat{d}_q x_q^s$ gives the mass of users of class q in the cell s. For ease of notation, we define $y_q^s \triangleq \hat{d}_q x_q^s$. Thus, the total mass of users of class q is just $\sum_{i=1}^{S_q} y_q^i = \hat{d}_q$. Also define $w_q^s \triangleq \frac{L^s}{R_q^s}$. Under this notation, the throughput *per unit mass* is given by

$$T_q^s(\mathbf{y}^s) \triangleq \frac{L_q}{\kappa \sum_{j=1}^{Q} y_j^s + \sum_{j=1}^{Q} y_j^s w_j^s}. \tag{9}$$

In the above expression, L_q is the frame size in bits for users of type q. The denominator is the total time in seconds that the user has to spend in the system in order to successfully transmit these L_q bits. The ratio thus yields the throughput in bits per second.

Costs and Payoffs

We now consider the costs and payoffs in the system, which will all be measured in units of throughput. The total system throughput is given by

$$\mathcal{T}(\mathbf{y}) \triangleq \sum_{k=1}^{S} \sum_{i=1}^{Q} \tau_i^k(\mathbf{y}) = \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^k T_i^k(\mathbf{y}). \tag{10}$$

We consider this total to be the cost borne by the ISP. We assume that the ISP would like to maximize the system throughput, but would like to recover the cost, i.e., it is *individually rational*. Now, all the users in a cell should not be charged the same amount even if their throughput happens to be the same. We choose discriminatory pricing for the following reason. Given a time interval (even if the throughput is identical), there are some users taking only a small time share and others who take a large time share. The time share that a user occupies depends on the PHY rate and the frame size that he or she uses – clearly, one has to charge more for those who occupy a larger time share. This "occupancy factor" per unit mass is given by

$$\delta_q^s(\mathbf{y}) \triangleq \frac{\kappa + w_q^s}{\kappa \sum_{j=1}^{Q} y_j^s + \sum_{j=1}^{Q} y_j^s w_j^s}. \tag{11}$$

The occupancy of all users of a class q in cell s is $\delta_q^s(\mathbf{y}) y_q^s$. It gives the ratio of time occupied by users of class q to the total amount of time used by all users. Thus, a lower occupancy means that a class is being more efficient. Hence classes which have a greater value of occupancy ought to be charged more than those with a lower one. Since we measure payoffs in throughput units, we need to convert occupancy to throughput. In terms of throughput, the cost of supporting users of class q in cell s, in terms of the effect on throughput is the occupancy times the total throughput of all users in the cell. Thus, from the ISP's perspective, the cost of a unit mass of users of class q is

$$C_q^s(\mathbf{y}) \triangleq \delta_q^s(\mathbf{y}) \sum_{i=1}^{Q} \tau_i^s(\mathbf{y}) \tag{12}$$

In effect, the cost is "proportionally fair" – the more you occupy the more you must pay.

Now, a user would like to get as many frames of data in the time that he or she spends in the system. Clearly, users would like to maximize their individual throughputs for the price paid so the population would split up in such a way that this selfish objective is achieved. The payoff function per unit mass for users of class q in cell s is

$$F_q^s(\mathbf{y}) \triangleq T_q^s(\mathbf{y}) - C_q^s(\mathbf{y}). \tag{13}$$

The above expression tells a user the value of associating to a particular AP. The vector $\mathbf{y}$ is the strategy profile of all the users, which may also be considered as the state of the system. The strategy of users of a particular class q is the vector $[y_q^1, y_q^2, ..., y_q^{S_q}]$. Users would vary their strategies with time based on the state of the system in a manner that would give them the maximum payoff.

We illustrate the fact that the price being charged is actually the "cost-price". By this we mean that the total revenue obtained in a cell is identical to the total throughput in the cell (revenue is measured in the units of throughput). Consider the total revenue generated in a cell, obtained from (12), which is given by

$$\begin{aligned}
\sum_{i=1}^{Q} C_i^s(\mathbf{y}) y_i^s &= \sum_{i=1}^{Q} \left(\delta_q^s(\mathbf{y}) \sum_{j=1}^{Q} \tau_j^s(\mathbf{y}) y_i^s \right) \\
&= \sum_{j=1}^{Q} \tau_j^s(\mathbf{y}) \sum_{i=1}^{Q} \delta_q^s(\mathbf{y}) y_i^s \\
&= \sum_{j=1}^{Q} \tau_j^s(\mathbf{y}) \sum_{i=1}^{Q} \frac{(\kappa + w_q^s) y_i^s}{\kappa \sum_{j=1}^{Q} y_j^s + \sum_{j=1}^{Q} y_j^s w_j^s} \\
&= \sum_{j=1}^{Q} \tau_j^s(\mathbf{y}),
\end{aligned}$$

which is the total throughput in the cell. We have used the definition of occupancy (11) in the above derivation. Since the revenue is the same as the throughput in the cell, we have assumed that service is provided to just try and break even. The objective is to maximize the total system throughput.

It is clear that there is an inherent tussle between the users who are interested only in their individual payoffs and the global objective of trying to maximize efficiency of the system. The price of allowing users to multihome is the cost that is borne by the system. We will study this problem in detail in the next two sections.

V. The Non-cooperative Multihoming Problem

As mentioned in the previous section, users behave selfishly with each user trying to maximize his or her individual payoff by multihoming. Thus, we have a system where populations partition themselves among the different actions available to them. Hence the scenario fits into the paradigm of population games with Q classes of users. We denote the non-cooperative game by F. We would like to know how this system of competing users evolves in time. Would it converge to any particular state? In order to answer this question, we need to assume something about the dynamics of the users. As explained in section II, we model user behavior using dynamics of replicator or BNN type. We find a potential function , which can be used to convert the population game to the potential game framework. We show below that the total system throughput is a potential function for the game.

Theorem 3: The function

$$\mathcal{T}(\mathbf{y}) = \sum_{k=1}^{S}\sum_{i=1}^{Q} y_i^k T_i^k(\mathbf{y}). \tag{14}$$

where $y_i^j = 0$ if AP j is not available to user i is a potential function for the game F.

Proof: We have

$$\begin{aligned}
\frac{\partial \mathcal{T}(\mathbf{y})}{\partial y_q^s} &= \frac{\partial}{\partial y_q^s}\sum_{k=1}^{S}\sum_{i=1}^{Q} y_i^k T_i^k(\mathbf{y}) \\
&= \frac{\partial}{\partial y_q^s}\sum_{k=1}^{S}\sum_{i=1}^{Q} \frac{y_i^k L_i}{\kappa\sum_{j=1}^{Q} y_j^k + \sum_{j=1}^{Q} y_j^k w_j^k} \\
&= \frac{L_q}{\kappa\sum_{j=1}^{Q} y_j^k + \sum_{j=1}^{Q} y_j^k w_j^k} \\
&\quad - \frac{(\kappa + w_q^s)\sum_{i=1}^{Q} y_i^s L_i^s}{\left(\kappa\sum_{j=1}^{Q} y_j^k + \sum_{j=1}^{Q} y_j^k w_j^k\right)^2} \\
&= T_q^s(\mathbf{y}) - \delta_q^s(\mathbf{y})\sum_{i=1}^{Q}\tau_i^s(\mathbf{y}) \\
&= F_q^s(\mathbf{y}),
\end{aligned}$$

which means that $\mathcal{T}(\mathbf{y})$ satisfies the definition of a potential function. ∎

We have thus shown that the conditions required for Result 2 to hold are satisfied, and hence the stationary points of the dynamics are asymptotically stable. From the form of the potential function, we expect that the stationary point actually maximizes the throughput. In the following section we provide a characterization of the stationary point and show that this is indeed true. The potential function is non-negative, and the strategy space $\mathcal{Y}$ is a compact set. However, it is not a concave function. Hence, the potential function could have non-unique maxima.

We illustrate the fact that non-uniqueness of Wardrop equilibria are reflected in non-unique maxima of the potential function in the following example.

Example

Consider the simple case where there is only one class of users. From (13), we have that the payoff to a unit mass of users is $T_q^s(\mathbf{y}) - T_q^s(\mathbf{y}) = 0$ in all cells s. This means that all strategy vectors $\mathbf{y}$ yield equal payoffs, i.e., any state is a Wardrop equilibrium. The potential function is merely $\sum_{k=1}^{S_q}\frac{L_q}{\kappa+w_q^k}$ regardless of $\mathbf{y}$. Thus, in this case the potential function is maximum for all states of the system, which is consistent with the above. ∎

VI. The Price of Anarchy

Consider the dynamics of the previous section. We would like to know what the stationary points of the system are and what it means for the system throughput. Essentially we would like to know what effect selfish multihoming has on the efficiency of the system. In most work on selfish routing (such as [15], [16]), it is found that the Wardrop equilibrium is inefficient, i.e., system performance suffers in some way because of users being allowed to take selfish decisions. This inefficiency is referred to as *the price of anarchy*. In [17] it is shown that under fairly strong symmetry conditions on the payoffs and potential function, efficiency can be achieved. Our scenario does not admit such strong conditions due to the structure of the system considered. However, we have just seen that the total throughput acts as a potential function for the system, from which we expect that our pricing mechanism is efficient. But we have not characterized the stationary point to see what it actually looks like. Below we provide a characterization of the stationary point and show that it is efficient.

For both the replicator (1) as well as the BNN dynamics (3) , we have that $\dot{y}_q^s = 0$, implies that either

$$\begin{aligned}
F_q^s(\hat{\mathbf{y}}) &= \frac{1}{\hat{d}_q}\sum_{i=1}^{S_q}\hat{y}_q^i F_q^i(\hat{\mathbf{y}}) \\
&\text{or} \\
\hat{y}_q^s &= 0,
\end{aligned} \tag{15}$$

where we use $\hat{\mathbf{y}}$ to denote a stationary point. The above relations mean that users of class q would get identical payoffs in all APs that they use at equilibrium.

Now, consider the stationary point again. We define $\hat{F}_q \triangleq \frac{1}{\hat{d}_q}\sum_{i=1}^{S_q}\hat{y}_q^i F_q^i(\hat{\mathbf{y}})$. Then the stationary point conditions look like Kuhn-Tucker first order conditions of an optimization problem. Let us identify the Lagrange dual function associated with the above expressions. It is seen that the minimization problem

$$\min_{\lambda}\max_{\mathbf{y}}\left(\sum_{k=1}^{S}\sum_{i=1}^{Q} y_i^k T_i^k(\mathbf{y}) - \sum_{i=1}^{Q}\lambda_i\left(\sum_{j=1}^{S_i} y_i^j - \hat{d}_i\right)\right), \tag{16}$$

yields (15) as the Kuhn-Tucker first order conditions with $\hat{F}_i = \lambda_i\ \forall i \in \{1, 2, ..., Q\}$ We then have the following theorem:

Theorem 4: The equilibrium of the non-cooperative game F is identical to the solution of the constrained optimization problem

$$\max_{\mathbf{y}} \left(\sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^k T_i^k(\mathbf{y}) \right) \quad (17)$$

subject to the constraints

$$\sum_{j=1}^{S_i} y_i^j = \hat{d}_i \quad \forall i \in \{1, 2, ..., Q\} \quad (18)$$

and $y_i^j = 0$ if AP j is not available to users of class i.

Proof: From the above discussion we have that the non-cooperative game F converges to the solution of the Lagrange dual problem (16). Call the solution obtained as $\mathcal{T}(\hat{\mathbf{y}})$. Also, call the solution to the primal problem (17) as $\mathcal{T}(\mathbf{y}^\star)$. Now, the expression in (17) is not concave and there could exist multiple maxima. There could also be a duality gap between the primal and dual problem, i.e., $\mathcal{T}(\hat{\mathbf{y}}) \geq \mathcal{T}(\mathbf{y}^\star)$. But it is physically impossible for the system to converge to a state whose throughput is greater than the maximum possible, i.e., $\mathcal{T}(\hat{\mathbf{y}}) = \mathcal{T}(\mathbf{y}^\star)$ ■

As mentioned earlier, the set of stationary points contains the set of Wardrop equilibria and in the case of BNN dynamics, they are the same. In the case of replicator dynamics, the system state might either converge to a Wardrop equilibrium or get stuck at a boundary point.

The result which we have just seen, coupled with that of the previous section has interesting consequences. We have shown that multihoming users with dissimilar selfish dynamics being charged the cost price of their occupancy actually *optimize the system throughput*. In the language of the above literature, the result states that the price of anarchy using the pricing mechanism suggested is zero – anarchy is free!

The fundamental difference between our model and the work on selfish routing is that of multihoming – the fact that users do not need to choose a single AP, but can *split* traffic. The result that multihoming is efficient is somewhat reminiscent of a result in [13], which states that the stability region of the Internet is increased by allowing multi-path routing with traffic splitting at source using a suitable TCP version. In effect we say that "a little choice (selfish routing) may be bad but a lot of choice (multihoming) is good". It would be interesting to see if multihoming would perform efficiently on the Internet as a whole.

Simulation

We perform a simple experiment using Simulink to verify that selfish multihoming does indeed maximize the system throughput. In our simulation we assume that users use replicator dynamics. Consider the scenario where there are two classes of users. Both users have the same two APs available to them. Their values of frame size are identical and equal to unity. However, the values of $\kappa + L/C$ are different – class 1 users have parameters $[2, 1.5]$, while class 2 users have parameters $[1, 5]$ in the two APs. Under this scenario the throughput is maximized if all users of class 1 use AP 1 while all users of class 2 use AP 2, which is a degenerate case of multihoming. The throughput would then be 0.5 in AP 1 and 0.2 in AP 2. We illustrate that the throughputs do indeed converge to these values in Figures 2–3.

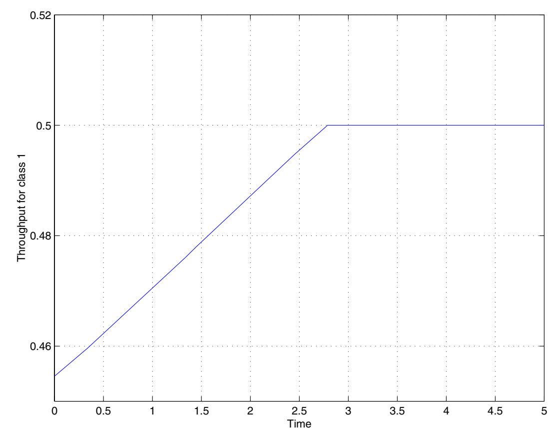

Fig. 2. Illustrating the convergence of the throughput of users of class 1.

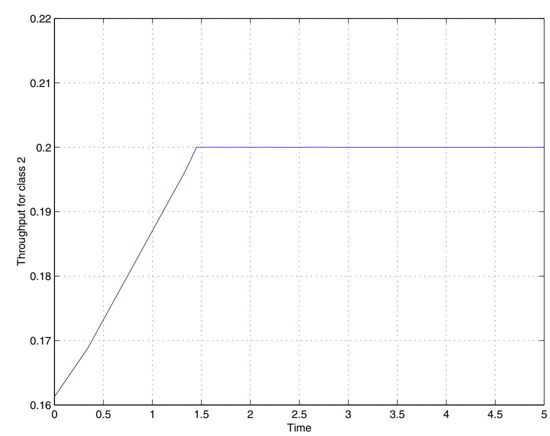

Fig. 3. Illustrating the convergence of the throughput of users of class 2.

Note that the case of all users of class 1 using AP 2 and all users of class 2 using AP 1 is a stationary point of replicator dynamics and is a boundary point of the state space, but is not a Wardrop equilibrium. However, any initial state except the one above would result in the state converging to the throughput optimal state.

VII. Economics of Multihoming

We have assumed so far that the ISP is a disinterested player and that the sole objective is to maximize the throughput of the system. However, this need not be the case in reality. We now consider a market model under which an ISP can charge more than the cost price for subscription. We make assumption that the potential mass of users in each class q is a fixed value denoted by $\hat{d}_q$. Clearly, even if prices were fixed, in practice one would expect a variation of user masses over the course of a day as they move around. However, we can think of this as the average that we can expect. The actual mass of users in the system would depend on the prices charged. Let the subscription price per unit mass charged in AP s to all users be denoted by P^s (independent of the PHY rate obtained). In addition to the subscription price, we assume that users are also charged the cost price of traffic described in the previous sections. We re-iterate that all prices are in units of throughput. The vector $\mathbf{P} \triangleq [P^1, P^2, ..., P^S]$ would determine

- the total mass of users in the system, and
- the way this mass gets partitioned between the different APs by multihoming.

To determine the total mass of users of each class, we need to make some assumptions about user demand. We assume that each class q is associated with a threshold value Λ_q. Users of class q would connect to an AP s if $P^s \leq \Lambda_q$. Once users connect to an AP, the throughput they obtain is determined by (9). The payoff per unit mass is then

$$F_q^s(\mathbf{y}) \triangleq T_q^s(\mathbf{y}) - C_q^s(\mathbf{y}) - P^s. \tag{19}$$

The ISP would like to maximize the profit regardless of the actual throughput of the system. The profit that the ISP makes is the difference between the total revenue and the cost (which we have assumed is equal to the actual throughput). Hence, the profit function of the ISP is merely

$$\rho_{multi}(\mathbf{P}) \triangleq \sum_{j=1}^{S} P^j \sum_{i=1}^{Q} y_i^j, \tag{20}$$

where $y_i^j = 0$ if users of class i do not connect to AP j.

The pricing scheme is somewhat similar to Paris Metro Pricing (PMP) [23]. In PMP a network is partitioned into several logically separate classes, with each having a fixed fraction of the entire network. Traffic in each fraction is handled using the same protocols, and no formal QoS guarantees are given to users. However, users in each fraction are charged different prices. The idea is that the higher priced fraction would be less loaded, thus leading to a higher perceived QoS. Like PMP, all users in a cell are given no QoS guarantee. If multiple APs are present in the same cell (perhaps owned by different ISPs), one could have "upper class" and "lower class" APs, which could charge different prices. However unlike PMP, our pricing scheme charges based on occupancy as well.

We showed in the previous section that multihoming along with a simple pricing mechanism maximizes the system throughput. We would like to know here whether the idea is economically feasible. If an ISP sets a price vector $\mathbf{P}$ for the APs in a region, would multihoming

- reduce or increase profit?
- always increase system throughput?

We now compare the profits obtained by the ISP and the throughput with and without multihoming and thus answer the question "What is the economic price of multihoming?".

Effect on Profit

To answer the question regarding ISP profit, we have to compare the profit when multihoming is an option and when it is not. So we need to know what users would do in the absence of multihoming. We make the assumption that users would connect only to (available) APs that display the lowest price. We denote this lowest price available to users of class i by $P_{min(i)}$. Then we have that the mass of users of class i connecting to AP k, under a given pricing vector $\mathbf{P}$ is such that $y_i^k = 0$ if $P^k \neq P_{min(i)}$. Under our assumption that the class as a whole follows the same dynamics, they would actually pick one of the APs displaying the lowest price. Thus, the profit function under unihoming is

$$\rho_{uni}(\mathbf{P}) \triangleq \sum_{i=1}^{Q} P_{min(i)} \hat{d}_i, \tag{21}$$

where $\hat{d}_i = 0$ if $P_{min(i)} > \Lambda_i$.

We are now ready to compare the two. We have the following theorem:

Theorem 5: For the same price vector $\mathbf{P}$, $\rho_{multi}(\mathbf{P}) \geq \rho_{uni}(\mathbf{P})$

Proof: The proof is straightforward once we realize that under a given price vector, the total user mass in the system is the same. Thus, we have

$$\begin{aligned}
\rho_{multi}(\mathbf{P}) &= \sum_{j=1}^{S} P^j \sum_{i=1}^{Q} y_i^j \\
&\geq \sum_{j=1}^{S} \sum_{i=1}^{Q} P_{min(i)}\ y_i^j \\
&= \sum_{i=1}^{Q} P_{min(i)} \sum_{j=1}^{S} y_i^j \\
&= \sum_{i=1}^{Q} P_{min(i)} \hat{d}_i \\
&= \rho_{uni}(\mathbf{P})
\end{aligned}$$

and we are done. ■

The result essentially says that there is no reason why an ISP should not allow users to multihome to its different APs. Any profit achievable when it allows unihoming can be met or exceeded by allowing multihoming.

Effect on Throughput

We now turn to the question of what effect multihoming has on the throughput of a system given a price vector $\mathbf{P}$. From the discussion of this paper so far, we would expect the throughput to be higher and here we show that this is indeed the case. We again have a game among the users. We would like to know what the equilibrium of the system would look like. As in the previous sections we identify a potential function for the system so as to convert it into a potential game.

Theorem 6: The function

$$\mathcal{T}_{multi}(\mathbf{y}) \triangleq \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^k \left(T_i^k(\mathbf{y}) - P^k\right). \tag{22}$$

where $y_i^j = 0$ if either AP j is not available to user i or $P^j > \Lambda_i$ is a potential function for the game F.

Proof: The proof is identical to that of Theorem 3 and is omitted. ■

As before we assume that when multihoming is an option, the population behavior is described by replicator dynamics, BNN dynamics or a combination of both. This would ensure that the stationary point of the system would be a Wardrop equilibrium (or a boundary value). As before, we have a primal-dual type of characterization of the stationary point. So we have the following theorem:

Theorem 7: The equilibrium of the non-cooperative game F is the solution of the constrained optimization problem

$$\max_{\mathbf{y}} \left(\sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^k \left(T_i^k(\mathbf{y}) - P^k \right) \right) \tag{23}$$

subject to the constraints

$$\sum_{j=1}^{S_i} y_i^j = \hat{d}_i \quad \forall i \in \{1, 2, ..., Q\} \tag{24}$$

and $y_i^j = 0$ if AP j is not available to users of class i or $P^j > \Lambda_i$.

Proof: Again, the proof is identical to Theorem 4. ■

We now assume that the users are not allowed to multihome. As mentioned earlier, they choose one of the APs displaying the lowest price. Let the AP that users of class q select be χ_q. Then the equivalent of $\mathcal{T}_{multi}$ is given as

$$\mathcal{T}_{uni} \triangleq \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^k \left(T_i^k(\mathbf{y}) - P^k \right), \tag{25}$$

where as usual,

$$y_i^k = \begin{cases} \hat{d}_i & \text{if } k = \chi_i \\ 0 & \text{otherwise} \end{cases}$$

Clearly $\mathcal{T}_{multi} \geq \mathcal{T}_{uni}$. We then have the following theorem on the throughputs in the two cases.

Theorem 8: Given a price vector $\mathbf{P}$, the system throughput when multihoming is permitted is at least that of when it is not.

Proof: Denote the equilibrium state when multihoming by $\hat{\mathbf{y}}$ and the state when unihoming by $\mathbf{y}^\star$. We have $\mathcal{T}_{multi} \geq \mathcal{T}_{uni}$

$$\Rightarrow \sum_{k=1}^{S} \sum_{i=1}^{Q} \hat{y}_i^k \left(T_i^k(\hat{\mathbf{y}}) - P^k \right) \geq \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^{k\star} \left(T_i^k(\mathbf{y}^\star) - P^k \right)$$

$$\Rightarrow \sum_{k=1}^{S} \sum_{i=1}^{Q} \hat{y}_i^k T_i^k(\hat{\mathbf{y}}) - \rho_{multi}(\mathbf{P}) \geq \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^{k\star} T_i^k(\mathbf{y}^\star) - \rho_{uni}(\mathbf{P})$$

$$\Rightarrow \sum_{k=1}^{S} \sum_{i=1}^{Q} \hat{y}_i^k T_i^k(\hat{\mathbf{y}}) \geq \sum_{k=1}^{S} \sum_{i=1}^{Q} y_i^{k\star} T_i^k(\mathbf{y}^\star),$$

since from Theorem 5 we have $\rho_{multi}(\mathbf{P}) \geq \rho_{uni}(\mathbf{P})$ (and they are non-negative). Hence the proof. ■

Thus, multihoming would do at least as well as unihoming in terms of throughput as well.

Optimizing the Profit

We have just seen that given any price vector, multihoming allows an ISP to increase both profit and efficiency. But how would this price vector be chosen? It is straightforward to show that the profit is always bounded as the user masses split within the compact set $\sum_{j=1}^{S_q} y_q^j \leq \hat{d}_q$. We show what the profit might look like for the simple case of a scalar price in Figure 4.

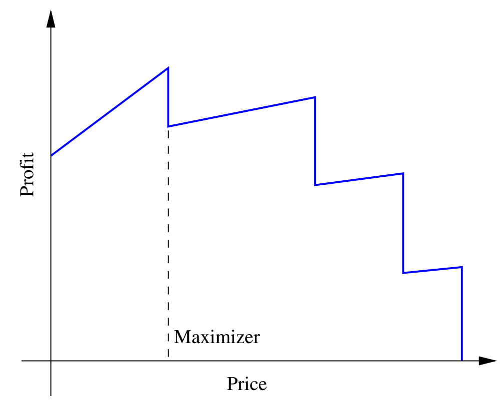

Fig. 4. Illustrating the fact that the total profit is bounded.

The vector $\mathbf{P}$ for which the profit is maximum need not be unique. Among all the maximizing $\mathbf{P}$ vectors, the ISP could choose the one that ensures the highest throughput of the system. Thus, economically there is a strong case for ISPs to allow multihoming to their APs.

VIII. Conclusion

In this paper we have sought to make a convincing case for ISPs to allow multihoming in IEEE 802.11 WLANs. We constructed a fluid model of user populations in a WLAN and understood how their throughputs varied with movement of user masses. We showed that users charged by a simple mechanism, using selfish dynamics would actually maximize the system throughput when allowed the option of multihoming. We thus established that under the multihoming scenario, anarchy comes at zero price. We also studied the economics of multihoming as seen by the ISP and showed that there is no loss of profit or throughput when users are allowed to multihome.

An important extension would be to devices associating with multiple technology access networks (WLAN and 3G for example) and could choose between technologies. In the future we would like to study the interaction of different ISPs, who might each own a different wireless LAN in the same region. Their interaction with each other and its effects on user throughputs would be of interest. We would also like to understand if results similar to what we have shown in WLANs applies to the Internet as a whole, i.e., can multihoming achieve efficiency on the Internet?

IX. Acknowledgment

The authors would like to thank the International Programs in Engineering (IPENG) at the University of Illinois at Urbana-Champaign and the Indo-French Center for Promotion of Advanced Research (IFCPAR, research contract No. 2900-IT) under whose funding this work was carried out.

References

[1] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance Anomaly of 802.11b," in *Proceedings of IEEE INFOCOM 2003*, San Fransisco, CA, March 2003.

[2] R. Chandra, P. Bahl, and P. Bahl, "MultiNet: Connecting to Multiple IEEE 802.11 Networks Using a Single Wireless Card," in *Proceedings of IEEE INFOCOM 2004*, Hong Kong, March 2004.

[3] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh, "A Comparison of Overlay Routing and Multihoming Route Control," in *Proceedings of the ACM SIGCOMM*, Portland, Oregon, February 2004.

[4] G. Tan and J. Guttag, "The 802.11 MAC Protocol Leads to Inefficient Equilibria," in *Proceedings of IEEE INFOCOM'05*, Miami, FL, March 2005.

[5] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on Selected Areas in Communications*, vol. 18, pp. 535–547, March 2000.

[6] A. Kumar, E. Altman, D. Miorandi, and M. Goyal, "New insights from a fixed point analysis of single cell IEEE 802.11 WLANs," in *Proceedings of IEEE INFOCOM 2005*, Miami, FL, March 2005.

[7] R. Venkatesh, A. Kumar, and E. Altman, "Fixed point analysis of single cell IEEE 802.11e WLANs: uniqueness, multistability and throughput differentiation," in *Proceedings of ACM SIGMETRICS 2005*, Banff, Canada, June 2005.

[8] Y. Bejerano, S.-J. Han, and L. E. Li, "Fairness and load balancing in wireless LANs using association control," in *Proceedings of ACM Mobicom*, Philadelphia, PA, September 2004.

[9] E. Altman, A. Kumar, D. Kumar, and R. Venkatesh, "Cooperative and non-coooperative control in IEEE 802.11 WLANs," in *Proceedings of the International Teletraffic Congress*, Beijing, August 2005.

[10] T. Korakis, O. Ercetin, S. Krishnamurthy, L. Tassiulas, and S. Tripathi, "Link Quality based Association Mechanism in IEEE 802.11h compliant Wireless LANs," in *Workshop on Resource Allocation in Wireless Networks (RAWNET)*, April 2005.

[11] A. Kumar and V. Kumar, "Optimal Association of Stations and APs in an IEEE 802.11 WLAN," in *Proceedings of of the National Conference on Communications (NCC)*, IIT Kharagpur, January 2005.

[12] H. Wang, H. Xie, L. Qiu, A. Silberschatz, and Y. R. Yang, "Optimal ISP subscription for Internet multihoming: Algorithm design and implication analysis," in *Proceedings of IEEE INFOCOM 2005*, Miami, FL, March 2005.

[13] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Overlay TCP for Multi-Path Routing and Congestion Control," to appear in *IEEE/ACM Transactions on Networking*.

[14] X. Lin and N. B. Shroff, "The Multipath Utility Maximization Problem," in *Proceedings of the 41st Allerton COnference on Communications, Control and Computing*, October 2003.

[15] T. Roughgarden and E. Tardos, "How bad is selfish routing?" in *IEEE Symposium on Foundations of Computer Science*, 2000, pp. 93–102.

[16] T. Roughgarden, "The price of anarchy is independent of the network topology," in *Proceedings of the 34th ACM Symposium on the Theory of Computing*, 2002, pp. 428–437.

[17] W. H. Sandholm, "Potential Games with Continuous Player Sets," *Journal of Economic Theory*, vol. 97, pp. 81–108, January 2001.

[18] H. Kameda, "How harmful the paradox can be in Braess/Cohen-Kelly-Jeffries networks," in *Proceedings of IEEE INFOCOM 2002*, New York, June 2002.

[19] A. Orda, N. Rom, and N. Shimkin, "Competitive routing in multi-user communication networks," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 614–627, 1993.

[20] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker, "On selfish routing in internet-like environments," in *Proceedings of the ACM SIGCOMM*, Karlsruhe, Germany, August 2003.

[21] J. G. Wardrop, "Some theoretical aspects of road traffic research," in *Proceedings of the Institute of Civil Engineers*, vol. 1, 1952, pp. 325–378.

[22] M. Patriksson, *The Traffic Assignment Problem: Models and Methods*. VSP Publishers, 1994.

[23] A. M. Odlyzko, "Paris metro pricing for the Internet," in *Proceedings of the ACM Conference on Electronic Commerce (EC'99)*, 1999, pp. 140–147.

[24] D. Ros and B. Tuffin, "A Mathematical Model of the Paris Metro Pricing scheme for charging packet networks," *Computer Networks*, vol. 46, no. 1, pp. 73–85, September 2004.

[25] R. E. Azouzi, E. Altman, , and L. Wynter, "Telecommunications Network Equilibrium with Price and Quality-of-Service Characteristics," in *Proceedings of the International Teletraffic Congress*, Berlin, September 2003.

[26] S. Shakkottai and R. Srikant, "Economics of Network Pricing With Multiple ISPs," in *Proceedings of IEEE INFOCOM 2005*, Miami, FL, March 2005.

[27] B. Hajek and G. Gopal, "Do greedy autonomous systems make for a sensible Internet?" in *Conference on Stochastic Networks*, 2002.

[28] J. Hofbauer and K. Sigmund, *Evolutionary Games and Population Dynamics*. Cambridge University Press, 1998.

[29] G. W. Brown and J. von Neumann, "Solution of games by differential equations," *Contributions to the Theory of Games I, Annals of Mathematical Studies*, vol. 24, 1950.

[30] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991 (7th printing, 2000).

[31] A. Haurie and P. Marcotte, "On the relationship between Nash-Cournot and Wardrop equilibria," *Networks*, vol. 15, 1985.