Game Theoretic Modeling of WiFi Pricing

John Musacchio & Jean Walrand

Electrical Engineering and Computer Sciences
University of California
Berkeley, CA USA
{musacchj,wlr}@eecs.berkeley.edu

Abstract

In this work we study the relationship between a WLAN owner acting as a wireless access provider and a paying client. We model the interaction as a dynamic game in which the players have asymmetric information – the client knows more about her utility function than the access provider knows. We find that if a client has what we call a web browser utility function, it is a Nash equilibrium for the provider to charge the client a constant price per unit time, and that clients with sufficiently high valuations for the service pay the price. In contrast, we find that if a client has what we call a file transferor utility function, with a bounded file length, the client should be unwilling to pay until the final time slot of her file transfer. We also analyze a Bayesian model in which the provider does not know whether he faces a web browser or file transferor type client, and study the case where there is no bound on the client's file length.

1 Introduction

Deployment of Wireless LAN base stations is growing quickly. A large number of the base station deployments are by private individuals for use in their homes, and by private businesses for use among their employees. Many of these base stations could be used to allow Internet access to a much larger set of users, in particular, users from the general public that lie or are passing within communicating distance of the base station. However, owners of private WiFi networks often choose to encrypt their networks to prevent outsiders from accessing them. Without a mechanism for a potential client to compensate the owner of the network, the network owner has no reason to accept the increased network traffic and security risk that would come from allowing the public to access her network.

If it were possible to incentivize owners of existing private wireless base stations to open their networks to the public, as well as incentivize people and institutions to deploy base stations where there are gaps in coverage, the result might be nearly ubiquitous WiFi coverage. In contrast to cellular phone networks deployed by giant infrastructure providers, this ubiquitous access network would be deployed by thousands, perhaps millions, of autonomous self-interested agents.

The simplest way for a client to compensate a base station owner would be for the client to just pay the base station directly. We refer to this as a "direct" model. Other models are possible. One model that is becoming increasingly popular is the aggregator model. In the model, the deploying business partners with an aggregator, such as Boingo

[6]. The aggregator attaches its brand name to hot spots and ensures that a consistent product is offered among the hot spots deployed by different businesses. The aggregator also handles the billing for the service, and can offer the user subscription billing plans that apply to all of the branded hot spots. The aggregator collects the revenue from the client, and then redistributes some of the revenue to the deploying business partners. This model will likely continue to grow in popularity, but because it involves a third party mediator between base station and client it is not clear wether this model could scale to millions of base stations. For now, we will focus on studying the properties and viability of a "direct" model involving the two principle parties, base station and client, and minimal or no third party involvement.

1.1 Making the "direct" model viable

Though it is simple, a number of challenges exist in making a direct model viable. In many cases a client and base station may not know each other's identity, and may not be able to trust each other to carry out their side of a transaction. To understand the problem, imagine a scheme where the client pays for her entire session in one lump payment. In a scheme where the client pays the base station in advance, or pre-pay scheme, a malicious base station might accept payment and then fail to deliver service. In a post-pay scheme, a client may fail to make a promised payment after receiving service. In fact, if we imagine that the base station and client are players in a game, and are trying to maximize their reward from this single transaction, a client *should* try to obtain service without paying, and a base station *should* try to take the client's money without giving him any service. Therefore, we must take care in structuring the game in a way that deters the players from cheating. Yet at the same time, we would like to avoid introducing a third party enforcement agent to the game. In an implementation, an enforcement agent would probably have to be centralized, and thus might limit scalability.

One variant of a direct payment model would be for the client to pay the base station in small amounts over the duration of the session. The intuition justifying this scheme is that a base station will want to play "fair", lest it be punished by being denied payments in the future, and a client will want to keep paying throughout the session to ensure that its service is not cutoff. We must also be concerned with how the base station changes its access price over the duration of a session. Will the base station entice the client to connect with a low price in the beginning, and then later threaten to cutoff the client's file transfer unless she agrees to pay a new higher price rate? Will the client refuse to connect to a base station out of fear that the price will be unstable during the duration of a session? These are the kinds of question we hope to address in this work.

1.2 Overview

In Section 2 we introduce our basic two-player game model. In Section 2.1 we discuss an instance of the basic model that we call the web browsing model, and show that the provider or base station should charge his client a constant price. In Section 2.2 we introduce the file transferor model, and show that if the base station has an a priori bound on the possible length of the client's file, the client's dominant strategy is to refuse any price greater than zero until the final time slot of her file transfer. The base station, in turn, charges a price of zero, until a critical slot $\hat{i}$ is reached, and then tries to charge for the whole session in one shot. In Section 2.3, we study a model in which the base station is unsure wether he faces a file transferor or web browser client. In Section 2.4, we investigate what happens when a file transferor type client's file length has no a priori bound (a geometrically distributed file length for example.) Section 3 addresses the issue of finding an appropriate micropayment scheme for the direct payments model. Finally, Section 4 summarizes our results and discusses our ongoing work.

2 Basic Model

We can formulate the interaction between a base station and paying client using a simple two-player game model. The game progresses in discrete time slots or "periods." At the beginning of the first time slot, the base station proposes an access price, p_1 for access during the first time slot. The client can either accept the price and connect, or reject the price and not connect. If the price is rejected, the game ends and both client and base station receive zero payoff.

In general, the base station offers connectivity at the beginning of time slot i at price p_i. The game ends the first time the client rejects the base station's proposal.

On completion of the game, the client's utility function $f(t, K)$ is a function of the number of time slots the client was connected, and a parameter K which we call the client's intended session length. The client's net payoff is $f(t, K) - \sum_{i=1}^{t} p_i$. The base station's net payoff is simply $\sum_{i=1}^{t} p_i$. The underlying assumption is that the base station's marginal cost to provide the service to the client is negligible. We study the Nash equilibria of this game, under different assumptions of the structure of the utility function $f(t, K)$. We assume the reader has some knowledge of game theory in the discussion that follows. A reference for the subject is [3]. In particular, we make use of the concept of perfect Bayesian equilibrium. The concept of perfect Bayesian equilibrium (PBE) is an extension of the concept of subgame perfection, in cases where one or more players does not know the exact utility function of his opponents, but instead knows a probability distribution on a set of possible utility functions. A PBE is a strategy profile – or specification of each player's strategy – such that no player can increase her expected payoff by unilaterally deviating from her PBE strategy at any point in the game.

2.1 Web Browsing Model

In a version of the basic model we call the web browsing model, the client's utility is proportional to the length of time t connected, up until the maximum intended session length K.

$$f(t, K) = U \cdot \min(t, K) \tag{1}$$

The parameters U and K have a random distribution. The client knows the sample value, the base station just knows the distribution.

Theorem 1 *Consider a web browser client with utility defined by expression (1). Suppose that U and K are independent and finite mean. Then the following strategy profile is a PBE equilibrium.*

- *The client connects or remains connected in each slot iff $i \leq K$ and $p_i \leq U$. (We refer to this as the "myopic strategy.")*

- *The base station charges price p^* in all time slots, where p^* is a maximizer of $xP(U \geq x)$.*

Proof: First, we find the base station's optimal counter strategy to a client playing the "myopic strategy." A pure strategy for the base station can be specified by a sequence of prices $\{p_i\}_{i=1..\infty}$ to charge at each time slot i. After a client accepts a price p_i in slot i, the base station knows that he can charge price p_i in slot $i+1$ without risking the client leaving. Therefore, the base station's best strategy against a myopic client will be a non decreasing price sequence. Call the set of non decreasing price sequences P^+. The base station wishes to maximize

$$\max_{\{p_i\}_{i=1..\infty}\in P^+}\left[\sum_{i=1}^{\infty} p_i P(p_i < U)P(K \geq i)\right]. \tag{2}$$

Because U and K are finite mean, one can substitute their markov bounds into expression (2) to show that the base station's expected payoff against a myopic client is bounded. Because the sequence where $p_i = p^* \ \forall i$ maximizes each term of expression (2), and because this sequence is an element of P^+, it is an optimal sequence of prices for the base station to charge a myopic client.

Now looking at the client's side, it is easy to see that the myopic strategy is a best response to a base station that never lowers prices.

Because these strategies are best responses to each other, and because one can repeat the above arguments to find that the same strategies are best responses to each other in any continuation game after any slot $i \ \forall i$, the strategy profile is a PBE. ■

It is somewhat surprising that the base station would keep its price constant. As we said, whenever a myopic client accepts price p_i, the base station can refine its conditional distribution of U by lower bounding it by p_i. One might have expected that a base station might want to try charging a higher price than p_i after learning that the client's utility is at least p_i. We have shown here that this intuition is not correct.

We picked p^* to be a maximizing value of $xP(U \geq x)$. In cases where $xP(U \geq x)$ has a set of maximizers $\mathbf{X}$, it is a PBE for the base station to charge a prices such that $p_i \in \mathbf{X}$ and $p_i \geq p_j$ if $i > j$, and for the client to follow a myopic strategy.

Our web browser result is similar to results shown in other contexts in the economics literature. For example, in [5] the authors show that under certain assumptions that it is not more profitable for a seller to condition pricing on the past behavior of the customer.

2.2 File Transfer Model: Bounded Length

In an instance of the basic model that we call the file transfer model, the client's utility function is a step function. The client must remain connected for the entire intended session length, and complete his file transfer, to get any utility.

$$f(t,K) = \begin{cases} 0 & \text{if } t < K \\ K+\epsilon & \text{if } t = K \end{cases} \tag{3}$$

The ϵ of expression (3) is assumed to be positive and smaller than the smallest unit of payment. After the statement and proof of Theorem 2 we discuss the case when ϵ is set to 0.

Theorem 2 *Suppose the client has a file transfer utility function as in expression (3), the session length K, is distributed on $\{1,...,N\}$, with a sample value* known *to the client, and* unknown *to the base station. We also assume that the function $iP(K=i)$ has a unique maximizer, $\hat{i}$.*

Then the following strategy profile is the unique, perfect Bayesian equilibrium:

- *The client accepts the base station's offer if $p_i = 0$ when $i < K$, and $p_i \leq K$ when $i = K$. We refer to this as the "pessimistic" strategy.*
- *The base station charges:*

$$p_i = \begin{cases} 0 & \text{if } i < \hat{i} \\ i & \text{otherwise} \end{cases} \tag{4}$$

Proof: The proof uses backwards induction, and iterated deletion of dominated strategies. We begin by showing that clients with intended session length N (which we call type N clients,) follow the pessimistic strategy.

Type N Clients follow the pessimistic strategy:

Suppose the client has an intended session length, K, of N. We will refer to such a client as a type N client. When the game reaches slot N, a type N client's dominant strategy is to accept any price less than N, because doing so would earn her utility $N + \epsilon$ for finishing her file. When the game reaches slot N, the base station knows that he faces a type N client, and thus his dominant strategy, after deleting his client's dominated strategies, is to charge N. A type N client can predict that she will be charged N in slot N, and therefore she would be unwilling to pay positive prices in the previous time slots.(Accepting a positive price in a previous time slot would result in her finishing the game with a negative payoff, while quitting the game would leave her with a zero payoff.) Thus, a type N client plays the pessimistic strategy.

Suppose we have shown that Clients of Type X through type N follow the pessimistic strategy. Show for Type $X - 1$:

Suppose that a type $X - 1$ client does not play the pessimistic strategy by accepting a nonzero price in a slot with index less than $X - 1$. When the game reaches slot $X - 1$, the base station can deduce that the client is of type $X - 1$ or greater, and that clients of type X or greater would have already quit the game. Thus, the base station knows he faces a type $X - 1$ client. Knowing his client's file ends in slot $X - 1$, he can charge $X - 1$ and be assured that the client will be compelled to pay. This outcome should have been predictable to the type $X - 1$ client, thus she would not have been willing to pay a nonzero price in a slot of index less than $X - 1$. We have shown that clients of type $X - 1$ play a pessimistic strategy.

By induction, clients of all types play the pessimistic strategy.

Base Station counter strategy:

A base station facing pessimistic clients has only one chance to charge nonzero prices. The base station charges according to expression (4) and choose $\hat{i}$ to maximize $iP(K = i)$. ∎

When $iP(K = i)$ has more than one maximizing value, one can see that clients still play the pessimistic strategy, but the base station will have more than one optimal best response to pessimistic clients. In particular, the base station can pick any maximizing index of $iP(K = i)$ as the slot index to start charging nonzero prices.

The proof of Theorem 2 shows that for any positive ϵ, we can find a unique strategy profile by iterated deletion of dominated strategies. When ϵ is set to zero, the strategy profile described in the statement of Theorem 2 is still a perfect Bayesian equilibrium,

p	Base Station Slot 1 Price	Base Station Slot 2 Price	WB Client Strategy	FT Client Strategy
$0 < p < \frac{1}{2}$	1	1	myopic	myopic
$\frac{1}{2} \leq p \leq \frac{2}{3}$	1	1 w.p. $1 - 2\epsilon$ 2 w.p. 2ϵ	myopic	myopic w.p. $\frac{1-p}{p}$ pess. w.p. $\frac{2p-1}{p}$
(2 equilibria)	0	2	myopic	pessimistic
$\frac{2}{3} \leq p \leq 1$	0	2	myopic	pessimistic

Table 1: The perfect Bayesian equilibria of the fixed length Bayesian game. The "myopic" strategy for clients of both types, is to accept a price of 1 or less in both slots. The "pessimistic strategy" for file transferor (FT) clients is to accept no price larger than zero in the first slot, and no price larger than 2 in the second slot.

but technically it is no longer unique, because the client's best responses are no longer unique. For example, a type 1 client facing a price of 1 in slot 1 is indifferent between accepting or rejecting when $\epsilon = 0$.

2.3 Bayesian Game: Session length of 2.

Having looked at the web browser and file transferor cases, it is now interesting to consider a case where the base station does not know wether he is facing a web browser or a file transferor. We discuss a special case here where the intended session length is assumed to be 2.

Theorem 3 *Suppose that:*

- *The client is a file transferor (FT) with probability p and a web browser (WB) with probability $1 - p$.*
- *Clients of both types have a intended session length of 2, and this is known to all parties.*
- *Clients have a utility of $1 + \epsilon$ per time slot connected, but file transferors need to be connected in* both *slots to earn any utility. Web browsers do not have this requirement.*
- *ϵ is smaller than the smallest unit of payment, and $\epsilon > 0$.*

Then the perfect Bayesian equilibria of this game are determined by the parameter p as described in Table 1.

The proof of Theorem 3 is provided in the appendix.

It is encouraging that for $p < \frac{1}{2}$, the only PBE has the base station charging a constant price. This suggests that if the base station thinks that its client is more likely to be a web browser than a file transferor, the base station's will behave as it does in the pure web browsing model of Section 2.1. Also when $p < \frac{1}{2}$, both file transferor and web browser clients accept the price of 1 in both time slows. When $\frac{1}{2} \leq p \leq \frac{2}{3}$, it becomes much harder to predict how the game will play out. There are two PBE, and furthermore, one of the PBE is a mixed equilibrium. When $\frac{2}{3} \leq p$, the game again becomes more

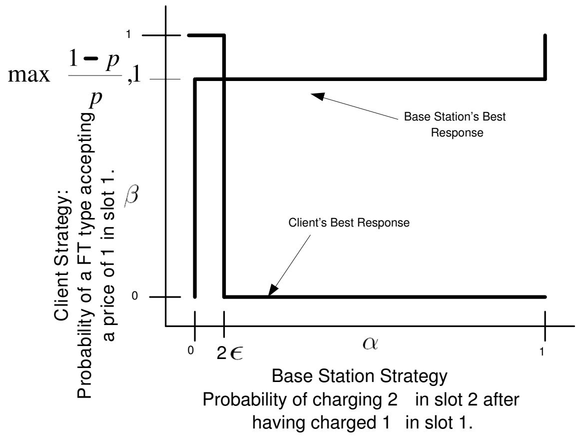

Figure 1: Analysis of the fixed session length bayesian game: Best responses in the continuation game following the base station charging 1.

predictable, and there is a single PBE. FT type clients play a pessimistic strategy, and the base station charges $\{0, 2\}$.

When the ϵ of Theorem 3 is set to zero, the strategy profiles given in the theorem statement of Theorem 3 are still perfect Bayesian equilibria. However, the best response sets of the clients grow because they are indifferent about accepting prices that are exactly equal to their utilities. Also, the mixed equilibrium in the $\frac{1}{2} < p < \frac{2}{3}$ changes in character significantly, which can be seen by studying the right plot of Figure 2.3. When $\epsilon = 0$, the best response curves in the continuation game after the base station charges 1 intersect a long a line segment where the base station charges 1 in slot 2 with probability 1, and the FT clients accept a price of 1 in the first slot with any probability between 0 and $\frac{1-p}{p}$.

2.4 Bayesian Game: Unbounded Length

The analysis of the finite horizon file transfer model depended heavily on having a fixed end point from which to work back. One might think that in an infinite case, where there is at each step some probability of the game continuing another step, that the solution would be very different. One might hope that the base station offering a constant price could be part of a PBE. We show that this is not the case. A constant price strategy cannot be part of a PBE.

Theorem 4 *Assume that*

- *The intended session length* K *is distributed on*$\{1, 2, ...\}$.
- *With probability* $p > 0$ *the client has a file transferor utility function:* $f(t, K) = K \cdot 1(t \geq K)$.
- *With probability* $1 - p$, *the client has a web browser utility function:* $f(t, K) = \min(t, K)$.
- $P(K = i | K \geq i) > \delta$ *for some* $\delta > 0$
- $E[K = i | K \geq i]$ *finite.*

Then, constant price is not a perfect Bayesian equilibrium.

Proof: Suppose the game has reached time slot i, and that the base station has been charging a constant price of $x : 0 < x < 1$ in slots $\{1, ..., i-1\}$. The base station can deviate from constant price and charge $i - \epsilon$ in slot i. If the client is a web browser with intended session length exactly equal to i, the client would pay the $i - \epsilon$ price. Other types of clients would quit. The expected reward to the base station for this deviation is $p(i - \epsilon)P(K = i|K \geq i)$, while the expected reward for maintaining the constant price is $xE[(K - i)|K \geq i]$. One can always find a large enough i so that the expected reward for deviating to price $(i - \epsilon)$ exceeds the expected reward of sticking to constant price. ■

3 Payment Schemes

The model we have been analyzing assumes that the client pays the base station in small payments over the course of the session. Unfortunately, most electronic payment schemes in common use today have a relatively large transaction overhead. If the length of a time slot were only on the order of a few minutes or less, it is likely that the transaction overhead would be comparable in size to the size of the payment itself. Clearly, such a high payment overhead would be prohibitive. Much work has been done in the area of making small payments, or micropayments, with a minimum of overhead. One scheme in particular appears to be promising in this context. The PayWord scheme [2, 1] makes it possible for the payee, in our case the base station, to aggregate many payments from a client into a single, larger payment. The scheme works on the principle of pay chains. In the brief description that follows, we assume the reader is familiar with the ideas of digital signatures and one way functions. For a review of these concepts, the reader can consult [4].

Prior to beginning a session, a client would compute what Rivest the authors of [2] calls an "H-chain" by repeated evaluation of a one-way function H. More specifically, the "H-chain" consists of values $x_0, x_1, x_2, ..., x_n$ where $x_i = H(x_{i+1})$ for $i = 0, 1, ..., n-1$.

A client begins a session by passing the base station the root of the chain x_0 that has been signed with the client's digital signature. In each subsequent time slot, the client makes another payment by passing the next consecutive value of the H-chain to the base station.

When the base station is ready to deposit the micropayments, the base station can combine them into single deposit by passing to the bank the client's digitally signed x_0 and the value, x_i.

The bank can verify the authenticity of x_0 by its digital signature, and can verify that i micropayments were made by iterating H on x i times.

The most important property of this payment scheme, is that it does not involve a third party for each micropayment. The base station need only contact a certificate authority at the beginning of the session to verify the authenticity of the client's digital signature of x_0. The base station can independently verify the authenticity of successive micropayments by verifying $x_i = H(x_{i+1})$.

4 Conclusion

Our initial results show that a model in which the client pays the base station directly may be viable if the client's utility function is that described by expression (1), because the web browsing formulation leads to a constant price Nash equilibrium. However, we have shown that if the client has a file transfer utility function, we find that the base station charges for the session in one shot in a particular time $\hat{\imath}$. Clients who finish their file transfers before $\hat{\imath}$ get a free transfer, while those with files longer than $\hat{\imath}$ quit early. This is clearly not an efficient outcome. This inefficient outcome can be avoided in a mixed model where the file length is bounded, and the probability that the client is a file transferor is small, but it cannot be avoided if the length of the file has no priori bound.

We are currently studying other variations of our models. We are exploring models with combinations of the following features:

- Clients can make repeat visits to the same base stations.
- Multi-hop situations in which clients can serve as a base station to a third party.
- Dynamic arrival models of clients to base stations with capacity constraints.

5 Appendix

5.1 Bayesian Result

A base station would charge either 0 or 1 in slot 1. First, we examine the case where the base station charges 1.

5.1.1 Strategies in the continuation game after the base station charges 1 in slot 1

Suppose that a client of type FT believes that the base station will follow a price of 1 with a price of 2 with probability α. With this belief, a FT client connects in slot 1 if his expected payoff is positive, which can be reduced to the condition $\alpha \leq 2\epsilon$

Now suppose that a FT client accepts a price of 1 with probability β. When a client of unknown type accepts the base station's price in slot 1, the base station believes he is a facing a FT client with the conditional probability: $\beta p((1-p)+\beta p)^{-1}$ Knowing that only an FT client would accept a price of 2 in slot 2, the base station prefers charging 2 if

$$2\frac{\beta p}{(1-p)+\beta p} \geq 1. \tag{5}$$

This reduces to $\beta \geq \frac{1-p}{p}$, which is possible only if $p \geq \frac{1}{2}$.

When β is exactly equal to $\frac{1-p}{p}$, the base station has no preference between selecting a price of 1 or 2 in slot 2, or choose the mixed strategy of picking either price with some mixing probability. We can construct a mixed Nash equilibrium of the continuation game by overlaying the best response functions of the base station and FT clients, as illustrated in Figure 2.3. The mixed Nash equilibrium occurs when β, the probability a FT client accepts a price of 1 in slot 1, is $\frac{1-p}{p}$, and when α, the probability the base station follows a price of 1 with a price of 2, is 2ϵ. The Base Station's expected payoff under this strategy profile can be computed to be $4(1-p)$.

When p is less than $\frac{1}{2}$, the base station's best response in the continuation game is to charge 1 in slot 2, regardless of the client's strategy. Because the client knows this, client's of both types have a best response of accepting the base station's price of 1 in slot 1.

5.1.2 Strategies in the continuation game after the base station charges 0 in slot 1

When the base station charges 0 in the first slot, clients of both types choose to connect. In slot 2, the base station's best strategy is to charge 2 in slot 2 if $p > \frac{1}{2}$. Otherwise, the base station charges 1 in slot 2. The base station's payoff in this situation is $\max(1, 2p)$.

5.1.3 Strategies of the overall game

When $p < \frac{1}{2}$ the base station's dominant strategy is to charge the price sequence $\{1, 1\}$. The base station would surely earn less by charging $\{0, 2\}$, or by charging $\{1, 2\}$.

When $p \geq \frac{1}{2}$ the base station is "tempted" to follow a price of 1 with a price of 2 to exploit file transferors. The above analysis shows that there is a mixed equilibrium in the continuation game following the base station charging 1 in slot 1. In this equilibrium, the base station and FT type clients play mixed strategies, and the base station's payoff in this equilibrium is $4(1-p)$. However, it is not a unique equilibrium. Another equilibrium is for FT clients to play a pure myopic strategy, and for the base station to charge $\{0, 2\}$. The base station's payoff is $2p$. Note that after the base station charges 0, the client's strategies in the continuation game are dominated. So by playing $\{0, 2\}$ the base station is assured a payoff of $2p$.

When $p > \frac{2}{3}$, it is better for the base station to play $\{0, 2\}$, and be assured of reward $2p$ than for the base station to charge 1 in the first slot and possibly earn $4(1-p)$ if both parties play the mixed equilibrium strategy profile in the continuation game that follows the charge of 1.

The equilibria strategies are summarized in Table 1.

References

[1] Silvio Micali and Ronald L. Rivest, "Micropayments revisited," in *CT-RSA*, 2002, pp. 149–163.

[2] Ronald L. Rivest and Adi Shamir, "Payword and micromint: Two simple micropayment schemes," in *Security Protocols Workshop*, 1996, pp. 69–87.

[3] D. Fudenberg and J. Tirole, *Game Theory.*, MIT Press, Cambridge, MA, 1991.

[4] J. Walrand, *Communication Networks: A First Course.* Aksen Associates Incorporated, Homewood, IL, 1991.

[5] A. Acquisti and H. R. Varian, "Conditioning Prices on Purchase History," Technical Report, School of Information Management and Systems, UC Berkeley, 2002.

[6] http://www.boingo.com.