

A Game-Theoretic Analysis of Wireless Access Point Selection by Mobile Users

Kimaya Mittal*, Elizabeth M. Belding†, Subhash Suri†

*Citrix Online, 6500 Hollister Avenue, Goleta, CA 93117

†Department of Computer Science, University of California, Santa Barbara

kimaya.mittal@citrix.com, {ebelding, suri}@cs.ucsb.edu

Abstract—A user located in a congested area of a wireless LAN may benefit by moving to a less-crowded area and using a less-loaded access point. This idea has gained attention from researchers in recent literature [3], [21]. However, its effectiveness and stability are questionable. Each user selects the access point that offers the optimal trade-off between load and distance to be traveled. Since users are selfish, a user’s selection may adversely impact other users, in turn motivating them to change their selections. Also, future user arrivals and exits may invalidate current selections. This paper presents the first game-theoretic analysis of this idea. We model access point selection as a game, characterize the Nash equilibria of the system and examine distributed myopic selections that naturally mimic selfish users. We analytically and empirically assess the impact of user diversity and dynamic exit patterns on system behavior. The paper contributes to a deeper understanding of the costs, benefits and stability of such a solution in various usage scenarios, which is an essential pre-requisite for real-world deployment.

I. INTRODUCTION

Wireless LANs are commonly deployed to provide Internet access, particularly in large areas such as university and corporate campuses, airports and shopping malls. To obtain network connectivity in a wireless LAN, a device must first locate and associate itself with an access point within transmission range. Several access points are typically deployed across an area in order to provide uninterrupted network coverage.

In many deployment scenarios, the distribution of users in the geographical area is non-uniform. Consider the example of an airport, where the concentration of users at different gates may vary significantly depending on flight schedules. A high concentration of users in a localized area increases the load on the access points in that area, resulting in a reduced bandwidth share per user. The wireless medium may become congested, resulting in undesirable effects such as packet loss and increased end-to-end delay.

While the access points in the congested area are overloaded, those in other less-crowded areas of the network may simultaneously be under-utilized. In such circumstances, a user located in the congested region could benefit by moving to a different area of the network and using an access point that is less loaded. This idea was first envisioned in a seminal paper by Satyanarayanan [21]. The paper describes a scenario where a user at an airport gate is unable to obtain sufficient bandwidth to complete her tasks due to a high concentration of users in the locality. As a solution to this problem, the smart computing system running on the user’s machine discovers a

nearby location with greater available capacity and prompts the user to move to that location. This idea is appealing since it not only meets the user’s needs but also redistributes users, balances network load and improves resource utilization.

Apart from being non-uniform, the distribution of users in the network may be dynamic and often unpredictable. This creates a potential problem, which we illustrate by extending the earlier example scenario. In the example, the user at the airport decides to move to a different gate in order to obtain more bandwidth. However, while she is positioned at the new location, more users arrive and cause an increase in the network load at this location. As a result, the bandwidth share available to the user decreases. The new bandwidth share could potentially be even lower than what was available to the user at the initial location. The user may decide to move again to receive more bandwidth. However, the same problem can recur at the new location. Thus, the change of location may bring no benefit to the user, resulting in dissatisfaction and wasted effort. In a worst case scenario, the user may be prompted to move repeatedly, leading to instability in the system.

Thus, although the idea of a mobile user changing location to obtain more bandwidth can bring benefits to both the user and the network, its effectiveness and stability need to be examined before it can be applied in real networks. With this motivation, we present an analysis of the stability and performance of such a system.

A user is motivated to move to a different location if and only if the bandwidth gain justifies the extra distance to be traveled. In other words, there is a trade-off between the bandwidth gained by changing location and the effort involved in traveling the additional distance to the access point. To select an access point, the user weighs the load and distance parameters for all access points in the system and then chooses the one that is optimal.

There are two main reasons for the potential ineffectiveness and instability of such a system. The first reason is that neither the network nor the user has prior knowledge of how the user distribution in the network will change in the future¹. As a result, decisions that are based on current conditions may be invalidated by future changes. The second reason is that each user makes an independent location change decision that optimizes her own benefit, without regard to other users

¹It may be possible to predict user movement based on past history and other domain-specific information, such as flight schedules for the airport example; however, this cannot be assumed for the general case.

or to the system as a whole. In other words, users are selfish entities and their decisions may therefore unintentionally harm the interests of other users and the entire system.

Game theory is an appropriate tool to analyze systems comprising selfish rational entities. We therefore model the access point selection system as a game. The stable user-to-access-point assignments are the Nash equilibria of the game.

Although game theory helps us explore multi-player interactions and understand the form of resulting stable outcomes, which can often be counter-intuitive, it does not suggest algorithms to reach the desired outcome. In deciding a strategy, each player must take into account the possible actions of other players; this potentially infinite sequence of speculation-counter-speculation often makes it difficult to design distributed algorithms for achieving Nash equilibria. Specifically, for our access point selection game, each user must consider the choices of other, both current and future, users. Since the choice of one user may invalidate the selections of other users, in principle there could be a non-terminating sequence of user moves, never leading to equilibrium.

Our analysis begins with a system comprised of a user population with homogeneous preferences and a simple arrival/departure model. Although our previous observations lead one to believe that an algorithm that guarantees equilibrium is likely to be fairly complex, we find that, with the simple system model, a trivial myopic algorithm that naturally mimics selfish user behavior always brings the game to a Nash equilibrium in just a single iteration. In other words, each user makes a single selfish selection based on the current state of the system and is never motivated to change her mind. This is different from several well-known games, where a greedy strategy does not result in equilibrium.

Real systems may have a heterogeneous user population, and users may arrive and depart from the system in a dynamic fashion. If the user population is heterogeneous, the myopic algorithm no longer results in equilibrium. However, we find that the maximum individual gain that a user may obtain through a unilateral change of strategy is bounded, and this bound is proportional to the extent of diversity. If diversity in the user population is limited, the potential individual gain is low, and therefore users have less motivation to change strategy, resulting in a greater likelihood that the system remains stable.

A dynamic user departure pattern is somewhat at odds with the concept of a Nash equilibrium since the definition of Nash equilibrium implicitly states that the user population must remain static. Dynamic departures can change the state of the system in unpredictable ways, and maintaining equilibrium may be impossible. We analyze the effectiveness of myopic access point selection based on a load-distance trade-off in such a system through simulations. Our results show that, under a realistic arrival/exit pattern, the effectiveness of myopic access point selection increases with the average number of users in the system. The experienced load of a user can be improved significantly by traveling relatively short distances. Further, the myopic access point selection of a user remains optimal for a large percentage of the time that the user is present in the system.

This paper thus contributes to a deeper understanding of the costs, benefits and stability of load-distance-based distributed access point selection in various usage scenarios. Such an understanding is an essential pre-requisite to enable real-world deployment of such a system. The remainder of the paper is organized as follows: Section II presents a brief review of related work. In Section III we model the access point selection game for a simple system, while Section IV describes the myopic access point selection algorithm. In Section V, we add some complexities to the system model and examine their effect on system behavior. Section VI presents our simulation-based evaluation of system performance with dynamic user arrivals and exits. Finally, Section VII concludes the paper.

II. RELATED WORK

The idea of a user changing location to obtain more bandwidth was first described by Satyanarayanan in the context of a pervasive computing environment [21]. In the envisioned scenario, a device gathers inputs from the network and application and appropriately suggests a location change to the user so that the user can obtain sufficient bandwidth to complete her tasks. However, no implementation or evaluation is performed and no consideration is given to the stability of the approach. Balachandran et al. propose network-directed roaming to relieve congestion in public area wireless LANs [3]. In their proposed solution, heavily-loaded access points direct mobile users to move to less-loaded cells for improved service. Since the new locations are computed centrally, stability is not an issue here; however, the drawback is that this scheme does not take into consideration the individual preferences of users. A solution that gives users the flexibility to make their own selections fits more naturally with the distributed nature of networks and the diversity and selfish nature of users. Sanzgiri et al. propose a decentralized protocol whereby users can change location to improve quality of service in wireless multihop networks [19].

Since the Internet is comprised of multiple players with selfish interests, game theory is an appropriate tool to study various Internet phenomena [15]. Specifically, game theory has been used to analyze network creation [6], [8], selfish routing [12], [16], [18] and transport protocols [1]. In the area of wireless networks, game-theoretic approaches have been proposed for several problems, including multihop routing [2], [24], power control [7], [20], participation incentives [22] and wireless access pricing [14], [10], [11]. In particular, in [10], [11] the fixed rate pricing model is shown to be invalid in limited capacity networks, and a fixed-rate, non-interrupted service model for wireless Internet access is proposed.

Suri et al. analyze the selfish load balancing game, where selfish clients make a server selection such that their individual execution latency is minimized [23]. This is similar to access point selection by selfish users. The authors ignore the network latency in their study, which is analogous to the distance between the user and access point in our work. Also, they impose additional restrictions on the set of permissible servers for a given client. Therefore, their results differ significantly from ours. In particular, the myopic strategy does not lead to a Nash equilibrium in their case.

III. THE ACCESS POINT SELECTION GAME

In this section we describe the wireless access point selection game. We begin with a detailed description of the system and state our assumptions. This is followed by a discussion of our game-theoretic model of the access point selection problem. We then derive the condition necessary for attaining a Nash equilibrium in this game.

A. System Description

The real-world problem of access point selection by mobile users involves several complexities. For example, users may have different bandwidth requirements and may assign different relative importance to the load and distance parameters. Users may enter and exit the system at random times. For our initial analysis, we ignore some of these complexities so that the problem becomes more tractable. A description of our system model, including all assumptions, follows. In Section V we revisit some of the complexities by introducing them to the system model and examining their effect on system behavior.

When a user enters the system (i.e. first connects to the network), she may be located anywhere in the geographical area covered by the network. A service is deployed in the network that informs users about the location and current load of all access points. On entering the system, a user obtains the necessary information from this service and then selects an access point after evaluating the load-distance trade-off. We assume that this evaluation and selection is done based on the user's requirements and preferences by software running on the user's machine and is therefore rational. The software then prompts the user to move to the selected access point.

A user can select any access point in the system. Selections are made sequentially, i.e. each user completes her access point selection before the next user enters the system. Since the access point selection process typically takes only a few seconds, the probability that another user arrives before the previous user completes her selection is low in most scenarios, and so we ignore this possibility in our analysis. Note that the user need not physically move to the new location before the next user's arrival; it is sufficient that the user's selection is communicated to the network service by the software running on the user's machine.

On selecting an access point, the user is assumed to move very close to the access point used, such that the distance between the user and the access point becomes negligible in comparison with the distances between access points. This assumption significantly simplifies our analysis. Once the user associates with an access point and begins a data session, she remains stationary. This is consistent with the observations of recent wireless network usage studies [9], [13].

For our initial analysis, we assume that all users exit the system within a very short time of each other. This is a reasonable assumption for users attending a sports/music event or a conference session, or users waiting for the same flight at an airport gate, etc. Our assumption implies that when the first user exits the system, all other users are approaching the end of their sessions and so are no longer motivated to change location. We therefore do not consider the effect of user exits

on the system equilibrium in our initial analysis. Dynamic user exits are considered in Section V-C.

All access points are identical in capability and differ only in location and traffic load. The transmissions of an access point and its associated users do not interfere with those of neighboring access points and their associated users; this is typically achieved by assigning non-overlapping frequency channels to neighboring access points.

Each user wants to maximize the bandwidth that she can obtain from the network and therefore prefers the least-loaded access point. It has been pointed out that access point load does not directly correlate with the number of associated users and depends more on individual workload behavior [4], [5]. However, user workloads may vary significantly during the course of their network association, thereby causing the access point loads to fluctuate. A selection based on fluctuating workloads is not likely to be stable. Also, it may be difficult for a user to pre-determine the workload she will generate. Therefore, for the access point selection decision, we believe that it is more practical to estimate access point load based on the number of associated users, which is a relatively stable property. Each user, therefore, prefers an access point that has fewer associated users. The effect of different individual workloads is considered in Section V-B.

Users prefer access points that are closer to their current location in order to reduce the extra distance they need to travel. In our initial analysis, we assume that each user assigns the same relative importance, or weights, to load and distance when selecting the optimal access point. In Section V-A, we examine the system behavior when users assign different weights to the load and distance parameters. All distances are Euclidean and therefore satisfy the triangle inequality.

In the following section, we describe how the access point selection problem is modeled as a game.

B. Selection Model

We model the access point selection game as follows. The players in this game are the users. Each user can select any one of the access points in the network. Thus, if there are M access points in the network, each player has M possible strategies. A cost is assigned to each strategy, i.e. each access point, as described further in this section. The utility of a player is inversely proportional to her cost. Each player's objective is to maximize her utility from the system by minimizing her cost. Therefore, every user selects the access point that has the least cost. Users make their selections sequentially. When making a selection, each user has knowledge of the decisions made by users before her in the sequence through information available from the deployed network service, but is unaware of the remaining users. The outcome of the game is an assignment of users to access points.

To determine how a cost may be assigned to each access point, we refer to our system description from Section III-A. We have assumed that a user prefers an access point that has fewer associated users. The user also wants to minimize the extra distance to be traveled in order to associate with an access point. Therefore, our cost function must have the following characteristics:

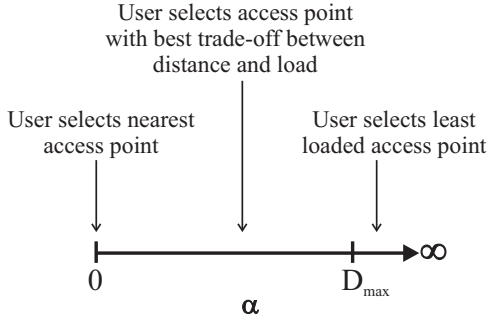


Fig. 1. Effect of the value of α on the access point selection.

- The cost should increase as the number of users associated with the access point increases.
- The cost should increase with an increase in the distance that must be traversed to use the access point.

There are several ways to define a cost function that satisfies the above characteristics. In this analysis, we choose a cost function that defines the cost assigned by user i to access point j as:

$$C_{i,j} = \alpha * x_j + D_{i,j} \quad (1)$$

where x_j is the number of users associated with access point j , $D_{i,j}$ is the cumulative distance traveled by user i to reach access point j , and α is a constant. This function satisfies both the desired characteristics of the cost function. Note that $D_{i,j}$ is the cumulative distance traveled by the user, i.e. if the user first selects access point j and then relocates to access point k , $D_{i,j}$ is the sum of the distances traveled for reaching access points j and k ².

The constant α expresses the relative weight assigned by the user to the load and distance parameters. Its value can vary from zero to infinity. The higher the value of α , the greater the weight assigned to the load parameter in calculating the cost of an access point. In other words, the higher the value of α , the greater the distance a user is willing to travel in order to use an access point with fewer associated users. Specifically, from equation 1 we can deduce that the user would be willing to travel a maximum extra distance of $n\alpha$ units in order to use an access point that has n fewer users associated with it.

If $\alpha = 0$, the user does not care about the access point load and always selects the access point that is nearest to her current location. As the value of α increases, a user is willing to travel larger distances in order to use a less-loaded access point. Let D_{max} be the *diameter* of the wireless network, i.e. D_{max} is the maximum distance between any two access points in the network. Then if $\alpha > D_{max}$, the user always prefers the access point with the least number of associated users without regard to its distance. For all values of α between 0 and D_{max} , the user is willing to trade distance to obtain a proportionate decrease in access point load.

The effect of the value of α on the access point selection decision is graphically represented in Figure 1. As specified in Section III-A, for our initial analysis we assume that all users

assign the same relative weight to the user load and distance parameters, so α has the same value across all users.

The access point selection game has some flavor of a sequential game: users make their decisions sequentially and the order of play can influence the outcome. Users are assumed to have perfect information about the selections of previous users. However, the selection of a user later in the sequence may cause previous users to change their decisions, and the number of moves may potentially be infinite. Further, at the time of making a selection, users have no knowledge about how many more users may enter the system and in what order. Therefore, the backward induction process used to solve sequential games cannot be used in this case. Also, since the selection by a player may cause multiple previous players to simultaneously change their strategy, the access point selection game is not a pure sequential game.

Next, we examine the conditions necessary to attain a Nash equilibrium in this game.

C. Nash Condition

A given assignment of users to access points is a Nash equilibrium if and only if no user can increase her utility, i.e. reduce her cost, by unilaterally changing her strategy. In other words, in a Nash equilibrium, no single user is motivated to move to a different access point if the selection of every other user remains unchanged. The only stable assignments are those that are Nash equilibria. In order to evaluate the stability of the assignments resulting from our algorithm, we first need to derive the condition that must be satisfied for an assignment to be a Nash equilibrium. This is the objective of this section.

Consider an assignment of users to access points. In this assignment, let user i be associated with access point j . Since we have assumed that a user moves very close to the access point used, user i must have traveled from her initial location to the location of access point j ; let this distance be $d_{i,j}$. Now, user i will be motivated to move to a different access point k if and only if the cost of access point k is less than that of access point j , i.e. $C_{i,k} < C_{i,j}$. If the assignment is a Nash equilibrium, no user is motivated to change her decision. Therefore, to attain Nash equilibrium, it is necessary that

$$C_{i,j} \leq C_{i,k} \quad (2)$$

$\forall i \leq N, k \leq M$, where N is the total number of users and M is the total number of access points. Note that j indicates the access point associated with user i in the given assignment. Using equation 1,

$$C_{i,j} = \alpha * x_j + d_{i,j} \quad (3)$$

where x_j is the total number of users associated with access point j , including user i , and $d_{i,j}$ is the distance of access point j from the initial location of user i (note that $D_{i,j}$ in equation 1 represents the cumulative distance).

Now, if user i moves to access point k , the number of users associated with access point k increases by one. Also, user i needs to travel an additional distance. Since user i is currently located very close to access point j , the current distance between user i and access point k is almost the same as that

²An alternative way to define the cost function is to consider the distance of the access point from the initial location of the user. The application scenario would determine which definition is more appropriate.

between access points j and k . So the additional distance to be traveled can be approximated as $d_{j,k}$, the distance between access points j and k . The cumulative extra distance traveled by user i in order to use access point k is then the sum of the distances traveled for reaching access point j and then moving to access point k , i.e. $d_{i,j} + d_{j,k}$. Therefore, the cost of moving to access point k is given by

$$C_{i,k} = \alpha * (x_k + 1) + (d_{i,j} + d_{j,k}) \quad (4)$$

where x_k is the number of users currently associated with access point k .

From equations 2, 3 and 4, we obtain the following condition for a Nash equilibrium:

$$x_j - x_k \leq \frac{1}{\alpha} * d_{j,k} + 1 \quad (5)$$

$\forall i \leq N, k \leq M$. Equation 5 represents the Nash condition for the access point selection game with our simple system model. Note that even though our cost function considers the cumulative distance traveled, the decision to move from one access point to another only depends upon the distance between the two access points without regard to the distance traveled previously. We use the Nash condition to evaluate the stability of assignments in further sections.

In the next section, we examine a simple distributed myopic algorithm for access point selection.

IV. DISTRIBUTED MYOPIC ALGORITHM

We noted in Section I that, although game theory enables us to understand the form of a stable outcome, it does not suggest a way to reach that outcome. In other words, the Nash condition, which we derived in Section III-C, does not provide an algorithm for attaining a Nash equilibrium. Algorithms that guarantee Nash equilibrium are often complex because of the inter-dependence of the actions of players and the speculation-counter-speculation involved.

We consider a simple myopic algorithm for access point selection and examine the stability of the resulting assignments. This algorithm naturally mimics the behavior of a selfish greedy user. We find that, under our simple system model, this algorithm always produces a Nash equilibrium in just a single iteration. This is unlike several other well-known games where a simple greedy strategy does not result in equilibrium.

A. Algorithm

On entering the system, each user executes the following algorithm to select an access point:

- 1) The user obtains information regarding the locations of access points and the number of users currently associated with each access point.
- 2) The user calculates the cost of each access point based on equation 1.
- 3) The user selects the access point that has the least cost.
- 4) The user associates herself with the selected access point after moving towards it if necessary.

The information service deployed in the network is critical in enabling this, or any other, access point selection algorithm.

Without this service, the user has no information about the location and current load of the access points in the system and therefore cannot make an informed decision. Note that, while making a selection, the objective of the user is simply to minimize her own cost.

The myopic algorithm naturally mimics the behavior of a selfish rational user in this system. With no knowledge about whether more users will enter the system and at what locations, the most logical strategy is to select the best access point based on the knowledge possessed, i.e. the current state of the system. Further, it is in the interest of the selfish user to directly associate with the preferred access point, thereby increasing its load and discouraging subsequent users from associating with it. The user gains nothing by temporarily associating with a different access point.

In the next section, we prove that this algorithm always produces a Nash equilibrium.

B. Proof of Nash Equilibrium

Consider an assignment of users to access points resulting from the myopic algorithm described in the previous section. Assume for the sake of contradiction that this assignment is not a Nash equilibrium. Then, at time T when all the users have completed their selection, there must exist a pair of access points j, k , for which the Nash condition from equation 5 is violated. In other words,

$$x_j(T) - x_k(T) > \frac{1}{\alpha} * d_{j,k} + 1 \quad (6)$$

where $x_j(T)$ and $x_k(T)$ are the number of users associated with access points j and k , respectively, at time T , i.e. in the final assignment. Note that all users are assumed to exit the system approximately simultaneously at some later time, so no user exits before time T .

Let user i be the last user to have selected access point j . We denote the time at which user i made her selection as t ($t \leq T$). Since user i preferred access point j over access point k at time t , it follows that

$$C_{i,j}(t) \leq C_{i,k}(t) \quad (7)$$

Using equation 1, we get

$$\alpha * x_j(t) + d_{i,j} \leq \alpha * (x_k(t) + 1) + d_{i,k} \quad (8)$$

where $x_j(t)$ and $x_k(t)$ denote the number of users at access points j and k , respectively, after user i has completed her selection (note that $x_j(t)$ includes user i) and $d_{i,j}$ and $d_{i,k}$ are the distances of access points j and k , respectively, from the initial location of user i . Therefore,

$$x_j(t) - x_k(t) \leq \frac{1}{\alpha} * (d_{i,k} - d_{i,j}) + 1 \quad (9)$$

Since user i is the last user to have selected access point j , the number of users associated with access point j cannot have increased after time t . In other words,

$$x_j(t) = x_j(T) \quad (10)$$

Also, the number of users associated with access point k cannot have decreased between times t and T since users do not exit the system before time T . So,

$$x_k(t) \leq x_k(T) \quad (11)$$

Subtracting equation 11 from equation 10, we get

$$x_j(t) - x_k(t) \geq x_j(T) - x_k(T) \quad (12)$$

Using equations 6, 9 and 12, we get

$$d_{i,k} > d_{i,j} + d_{j,k} \quad (13)$$

Since $d_{i,j}$, $d_{i,k}$ and $d_{j,k}$ are Euclidean distances, they must satisfy the triangle inequality, i.e.,

$$d_{i,k} \leq d_{i,j} + d_{j,k} \quad (14)$$

We see that equations 13 and 14 contradict each other. Therefore our assumption that the assignment resulting from the myopic algorithm is not a Nash equilibrium must be incorrect. We have thus proved by contradiction that every assignment of users to access points that is produced by the myopic algorithm is a Nash equilibrium with our simple system model. Observe that a Nash equilibrium is constantly maintained even while users are entering the system. In other words, all sub-assignments that exist temporarily while users are entering the system are also Nash equilibria.

In the next section, we add some complexities to our system model and examine how this affects the behavior of the system.

V. MORE COMPLEX SYSTEM MODELS

The analysis presented in the previous sections was based on the simple system model described in Section III-A. Real systems, however, are often more complex than this model. In this section, we examine how the behavior of the access point selection game changes when various complexities are added to the system model. First, we allow each user to assign different weights to the load and distance parameters in selecting the optimal access point. Next, we consider different individual user workloads when determining the access point loads. Finally, we discuss the effect of dynamic user exits from the system.

A. Diversity in α Values

The constant α represents the relative weight assigned by the user to the access point load and distance parameters. Different users may assign different values to α depending on their requirements. For example, bandwidth is of critical importance to a user who urgently needs to download a large file. This user may be willing to travel a relatively large distance in order to improve bandwidth availability; the value assigned by this user to α is therefore high. On the other hand, a user casually browsing the web does not care as much about bandwidth, so the value she assigns to α is likely to be low.

Let α_i denote the relative weight assigned by user i to the access point load and distance parameters. The new cost function is therefore

$$C_{i,j} = \alpha_i * x_j + D_{i,j} \quad (15)$$

The Nash condition from equation 5 correspondingly changes to

$$x_j - x_k \leq \frac{1}{\alpha_i} * d_{j,k} + 1 \quad (16)$$

\forall users i associated with access point j .

Notice that all users associated with a single access point j are no longer equivalent. Equation 16 indicates that, from among all users associated with a given access point, the Nash condition would be first violated for the user with the largest α value. In other words, the user with the largest α is the first to be motivated to change her strategy.

What event can motivate a user to change her strategy? We have assumed that all user exits happen approximately simultaneously and therefore do not motivate users to change their strategy. A user p associated with access point j can therefore be motivated to move only due to the arrival of another user q at the same access point. Since user q prefers access point j , it follows that $\alpha_q < \alpha_p$. In other words, a user is motivated to change her strategy only when another user with a smaller α value arrives at the same access point.

Is a Nash assignment always possible in this system? The arrival of each user can potentially displace only those users that have a larger α value than her own. These users may change their selection and in turn displace other users with still larger α values. Within a finite user population, this series of displacements cannot continue infinitely and must terminate when the user with the largest α value has made her selection. The resulting assignment is guaranteed to be a Nash equilibrium.

It follows that, with global knowledge of all users, a Nash assignment can be trivially computed by ordering the users in increasing order of α and then running the simple myopic algorithm from Section IV-A. However, when user arrivals are randomly ordered, the myopic algorithm is no longer guaranteed to produce a Nash assignment.

When an algorithm does not produce a Nash equilibrium, it is often useful to compute the maximum individual gain that a user can obtain through a unilateral change of strategy. The lower the individual gain, the lower the motivation for a change of strategy, and the higher the likelihood that the user adheres to her original selection, thereby maintaining the stability of the system. Consider a situation where a user q arrives at access point j and motivates another user p currently associated with access point j to move to access point k . Since user q prefers access point j over access point k , it follows that,

$$x_j - x_k \leq \frac{1}{\alpha_q} * d_{j,k} + 1 \quad (17)$$

From equation 15, the individual gain obtained by user p by moving from access point j to access point k is given by

$$\begin{aligned} C_{p,j} - C_{p,k} &= \alpha_p * x_j + D_{p,j} - (\alpha_p * (x_k + 1) + D_{p,k}) \\ &= \alpha_p * (x_j - x_k - 1) - d_{j,k} \end{aligned} \quad (18)$$

Using equation 17, we get

$$C_{p,j} - C_{p,k} \leq \left(\frac{\alpha_p}{\alpha_q} - 1\right) * d_{j,k} \quad (19)$$

Equation 19 gives the upper bound on the individual gain obtainable through a unilateral change of strategy. This upper bound depends on the maximum value of the ratio $\frac{\alpha_p}{\alpha_q}$. If the diversity of the user population is limited, the maximum value of this ratio is small, resulting in a small upper bound on individual gain. In other words, the less diverse the user population, the higher the likelihood of attaining a stable assignment through the myopic algorithm. In particular, if there is no diversity, i.e. all users have the same value for α , the maximum individual gain is zero and a Nash equilibrium is attained. This is in agreement with the result derived in Section IV-B.

B. Diversity in Individual Workloads

In our simple system model, we do not account for differences in individual user workloads, and we assume that the load at an access point is estimated based on the number of associated users. This assumption may not be accurate in some deployment scenarios. A user running a high bandwidth application, such as multimedia streaming or file download, is likely to increase access point load by a disproportionately higher amount than a user checking e-mail. Therefore, access point load may not directly correlate with the number of associated users.

Let s_i denote the individual workload of user i . The total load at an access point j is given by $S_j = \sum_i s_i \forall$ users i associated with access point j . Under this model, the original cost function from equation 1 changes to

$$C_{i,j} = \alpha * S_j + D_{i,j} \quad (20)$$

From this, the new Nash condition can be derived as

$$S_j - S_k \leq \frac{1}{\alpha} * d_{j,k} + s_i \quad (21)$$

\forall users i associated with access point j .

We again observe that all users associated with a single access point are not equivalent. Equation 21 indicates that the Nash condition at an access point j will be first violated for the user i with the smallest workload s_i . This can happen only when another user q arrives at the same access point (remember we have assumed that user exits do not motivate other users to move). Since user q prefers access point j , it must be true that $s_q > s_i$. Thus, in this system, the arrival of a user with a larger workload may displace another user with a smaller workload.

A Nash assignment is always possible in this system; this can be proved by an argument similar to the one used in Section V-A. With global knowledge, a Nash equilibrium can be computed by ordering users in decreasing order of individual workloads and running the myopic algorithm from Section IV-A. Since a user can only displace other users with smaller workloads, when users are ordered in decreasing order of workloads, no user can cause previous users to be displaced and a Nash equilibrium is obtained in a single iteration.

When users are randomly ordered, the myopic algorithm no longer produces a Nash equilibrium. We now compute the maximum individual gain that a user may obtain through a unilateral change of strategy. Consider that the arrival of a

user q at access point j motivates another user p currently associated with access point j to move to access point k . Since user q prefers access point j over access point k , it must be true that

$$S_j - S_k \leq \frac{1}{\alpha} * d_{j,k} + s_q \quad (22)$$

The individual gain obtained by user p by moving from access point j to access point k is given by

$$\begin{aligned} C_{p,j} - C_{p,k} &= \alpha * S_j + D_{p,j} - (\alpha * (S_k + s_p) + D_{p,k}) \\ &= \alpha * (S_j - S_k - s_p) - d_{j,k} \end{aligned} \quad (23)$$

From equations 22 and 23, we get

$$C_{p,j} - C_{p,k} \leq \alpha * (s_q - s_p) \quad (24)$$

This is the upper bound on the maximum individual gain, which depends on the extent of diversity in the user population. If diversity is limited, the value of $(s_q - s_p)$ is small, leading to a smaller individual gain and greater likelihood of stability. If we do not consider diversity among individual user workloads, i.e. $s_q = s_p$, no individual gain is possible and Nash equilibrium is attained, as derived in Section IV-B.

As mentioned in Section III-A, individual workloads in real deployments may vary significantly over time, thereby causing access point loads to fluctuate. So a decision based on current load conditions may potentially be invalidated in a short while. This leads to ineffective decisions and system instability. Also, a user may be unable to pre-determine her workload. Therefore, when deploying such a solution, it is more practical to use the number of associated users, which is a relatively stable property, as an indicator of access point load.

C. Dynamic User Exits

When a user exits the system, the load at the corresponding access point decreases, potentially motivating other users to change their previous selection and associate with that access point. In the analysis presented in the previous sections, we have assumed that all users exit the system within a short time of each other, and so the exit of one user does not motivate other users to change strategy since they are almost at the end of their own sessions. Although this assumption may be reasonable for certain deployment scenarios, there are several other scenarios where the exit pattern of users is more dynamic. Our objective in this section is to discuss the impact of dynamic user exits on the access point selection game.

As mentioned in Section I, a dynamic user exit pattern is somewhat at odds with the concept of a Nash equilibrium. The definition of a Nash equilibrium states that no single player is motivated to unilaterally change her strategy as long as all other players adhere to their decisions. This definition inherently implies that the player population must remain static. If users dynamically exit from the system, the user population may change in unpredictable ways. Under these circumstances, it is extremely difficult, if not impossible, to produce a user-to-access-point assignment that remains in equilibrium. For example, in a system consisting of two access

points and several users, no matter how the users are assigned to the access points, it is possible that all the users associated with one access point exit before the others, thereby disrupting the equilibrium. Further, in this scenario, no bound can be placed on the maximum individual gain obtainable through a change of strategy; as the number of users in the system increases, the individual gain can become infinitely large.

The above example is an extreme case. In real-world scenarios, it is unlikely that users exit the system in such a contrived fashion. Under a more realistic exit pattern, the potential individual gain through a change of strategy may not be as large. Further, as the system state changes in unpredictable ways with every user exit, users may be less likely to change their selections based on some intermediate system state. The question remains whether it is still beneficial for a user to base her access point selection on the load-distance trade-off when the user first enters the system, as opposed to simply associating with the nearest access point. In the next section, we examine this through simulation.

We use simulation to investigate the behavior of the system under dynamic exit patterns. The simulation results, presented in the following section, help us to gauge the benefits, costs and stability of access point selection based on the load-distance trade-off in dynamic environments.

VI. EVALUATION

In this section, we present a simulation-based study of myopic access point selection based on the load-distance trade-off. Our main objective is to evaluate the costs, benefits and stability of this approach with dynamic user arrivals and exits. We also study how the parameter α impacts the user-to-access-point assignments. The details of our experiments are presented in Section VI-A and metrics are defined in Section VI-B. In Section VI-C, we visually examine some user-to-access-point assignments and the corresponding load distributions. Section VI-D presents our simulation results.

A. Experiment Details

We evaluate the myopic access point selection under both uniform and non-uniform user distributions. Each scenario has 16 access points and 100 users placed in a 500mx500m area. The access points are arranged in a 4x4 grid in the center of the region, with 100m of distance between adjacent access points. For example, in Figure 2(d), the bounding box indicates the network area and the small squares represent the access points arranged in a grid. The user locations are generated randomly. In the first test scenario, users are uniformly distributed through the entire region, as shown in Figure 2(a), where each dot represents a user. In the second scenario, 80% of the users are placed within 16% of the total region (200mx200m area), as seen in Figure 2(d). This scenario models the phenomenon of high user concentration in a localized area, which can occur in real deployments as noted in Section I.

For our test scenarios, D_{max} , the maximum distance between any two access points, is approximately 425 units. We therefore vary the value of the parameter α from 0 to 500 in different tests. This covers the entire range of significant

α values indicated in Figure 1. We do not simulate diversity of α values and individual workloads among users in order to limit the number of simulation parameters and clearly identify the impact of dynamic exit patterns.

Simulating user arrivals and exits in a realistic manner is challenging. There are very few studies of real wireless network deployments that characterize user arrival and exit patterns [4]. Moreover, these patterns can vary significantly in different deployment scenarios. For our study, we assume that user arrivals follow a Poisson distribution; the Poisson arrival process is traditionally used to model independent user arrivals [17]. The time spent by users in the system, i.e. the *in-system time*, is assumed to follow an exponential distribution. Intuitively these distributions seem appropriate and provide a good starting point to examine the effect of dynamic arrivals and exits on the access point selection game. We leave the study of other distributions for future work.

The Poisson arrival process is characterized by a mean inter-arrival time (*MIAT*), while the exponentially-distributed time spent by users in the system is characterized by a mean in-system time (*MIST*). We fix the MIST at 300 seconds and vary the *MIAT* as 3, 10 and 30 seconds in different tests. In other words, the *MIST/MIAT* ratio is varied as 100, 30 and 10. The larger the value of this ratio, the higher the rate at which users arrive in the system compared to their in-system times, and so the greater the average number of users in the system. Apart from this dynamic arrival/exit model, we also implement the simple model described in Section III-A. In this model, users enter the system sequentially at the beginning of the simulation and exit simultaneously at the end.

In our simulations, when each user enters the system, she moves towards and associates with the access point that offers the best load-distance trade-off. Thereafter, further arrivals and exits may motivate the user to change her selection. However, since the arrivals and exits are dynamic and unpredictable, this motivation is likely to fluctuate. It is desirable for system stability that users adhere to their initial selections. We simulate this desirable behavior and measure metrics such as the potential individual gain that can be obtained by selecting a different access point and the time for which equilibrium is disrupted, i.e. at least one user in the system wishes to change her selection. These metrics indicate whether the simulated user behavior is effective and stable.

The simulations are implemented in Java. For each test scenario, we average our results over 100 simulation runs with the random number generator seeded differently in each run. In each run, the first user enters the system at time zero, and the simulation executes until the last user has exited the system.

B. Metrics

To evaluate the benefits, costs and stability of access point selection in a dynamic environment, we use the following performance metrics:

- **Maximum AP load difference:** We define access point (AP) load as the number of users associated with an access point. The maximum AP load difference is calculated as the difference in average load between the most-loaded

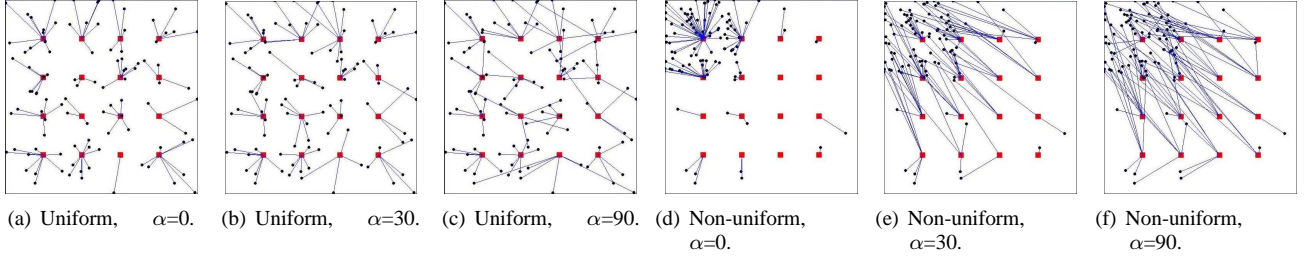


Fig. 2. Sample assignments with uniform and non-uniform user distribution for different values of α .




Fig. 3. Sample distribution of access point load with uniform and non-uniform user distribution for different values of α .

and least-loaded access points. A low value for this metric indicates a balanced load distribution.

- **Average experienced load per user:** The experienced load for a user is the load at the associated access point. This value is averaged over the in-system time of each user and then across all users. Again, a low value for this metric is desirable.
- **Distance traveled per user:** This is the distance that a user travels in order to use the access point with the best load-distance trade-off instead of the nearest access point. This metric indicates the cost of this approach.
- **System Nash time:** This is the fraction of time for which a Nash equilibrium exists in the system. The higher this value, the better the stability of the system.
- **Fraction of users motivated to change strategy:** This is the fraction of users that are motivated to move to a different access point at some instant during their in-system time. A low value for this metric is desirable for greater stability of the system.
- **Fraction of in-system time for which users are motivated to change strategy:** This is the fraction of a user's in-system time during which her initial access point selection is no longer optimal, thereby motivating her to move to a different access point. Only those users for whom this fraction is non-zero are included in the calculation of this metric. Again, a low value indicates greater system stability.
- **Maximum individual gain:** This is the maximum individual gain that a user can obtain by a change of strategy. It is calculated as a fraction of the user's current cost. The lower the value of this metric, the less the motivation for moving, and so the better the stability of the system.

C. Sample assignments and load distributions

Before studying the behavior of the system under dynamic user arrivals and exits, we visually examine some sample

user-to-access-point assignments and the corresponding access point load distributions. This gives us a feel for the effectiveness of this approach. Figure 2 shows sample assignments from test scenarios with both uniform and non-uniform user distributions. In each sub-figure, the small squares represent the access points, the dots represent the initial locations of users and the lines indicate the assignment. Figure 3 presents the final distribution of access point load corresponding to each assignment in Figure 2. Note that these assignments are with the simple system model without dynamic exits.

When $\alpha = 0$, there is no weight assigned to the access point load parameter, and so each user simply associates with the nearest access point as shown in Figures 2(a) and 2(d). There is considerable variation in the access point load for these assignments. As seen in Figures 3(a) and 3(d), respectively, the access point load varies from 1 to 13 users per access point under the uniform user distribution and from 0 to 44 users per access point under the non-uniform user distribution (note the difference in scale on the Y-axis in the two figures).

As the value of α increases, users travel larger distances to reduce load, and so we expect the access point load distribution to gradually even out. Figures 2(b) and 2(e) show the assignments when $\alpha = 30$. This is a relatively low α value when compared to the distance between adjacent access points (note that α represents the additional distance a user is willing to travel in order to associate with an access point that has one fewer associated user.) Even at this low value of α , the load distribution evens considerably when compared to the case where $\alpha = 0$. Figures 3(b) and 3(e) show the corresponding load distributions.

When the value of α increases to 90, the assignment lines intersect each other to a greater extent, as seen in Figures 2(c) and 2(f). This is because users travel larger distances on entering the system to select an access point with a lighter load. Note that the overall efficiency of the assignments could be improved by swapping user-access-point

pairs whose assignment lines intersect; this would maintain the same load distribution while reducing the total distance traveled. However, selfish users do not care about overall efficiency and only seek to optimize their individual load-distance trade-off. The load distribution improves significantly when α increases to 90 and is close to evenly balanced under both user distributions, as indicated in Figures 3(c) and 3(f).

We thus observe that myopic access point selection based on the load-distance trade-off is very effective in improving load distribution even at relatively low values of α . In the next section, we examine the effectiveness and stability when users arrive and exit in a dynamic fashion.

D. Experiment Results

We now present the results of our experiments. Due to lack of space, we only present the results from the test scenario with non-uniform user distribution. Results from the uniform user distribution scenario show similar trends. The graphs in Figure 4 show results for the dynamic arrival/exit model with different $MIST/MIAT$ ratio values, and also for the simple system model with simultaneous exits.

Figure 4(a) plots the maximum access point load difference for different values of α . At $\alpha = 0$, the maximum access point load difference is greater than 45 for the simultaneous exit model, indicating that load distribution is highly imbalanced. For this model, the load distribution improves dramatically with an increase in α , demonstrating effectiveness even under small values of α . As α continues to increase, load distribution evens and relative improvements become smaller in magnitude, until finally the distribution is perfectly balanced and no further improvement is possible.

With dynamic user arrivals and exits, the improvement is less dramatic. This is primarily because the average number of users in the system at any given time is less than in the simultaneous exit case. At $MIST/MIAT = 10$, users enter the system relatively slowly when compared to their average in-system time. As a result, the average number of users in the system at any given time is low, and so there is not much scope for improvement in load distribution. As the value of the $MIST/MIAT$ ratio increases, more users are present in the system at any given time, and so myopic access point selection is found to be more beneficial. The improvement in load distribution is most significant at low values of α . At higher α values, users are more evenly distributed and the average load of all access points over time is approximately equal. Therefore, the maximum difference between average AP loads is close to zero at higher values of α . (Note that the difference between instantaneous AP loads may still be high due to dynamic arrivals and exits).

The average load experienced per user, shown in Figure 4(b), shows similar trends for the same reasons. Figures 4(a) and 4(b) demonstrate the benefit of myopic access point selection in a dynamic system. The greater the average number of users in the system, the larger the benefit.

The distance traveled by users to associate with the optimal access point is shown in Figure 4(c). The distance increases with an increase in α since users attach greater importance

to access point load and are willing to travel a larger distance to associate with a lightly-loaded access point. As the $MIST/MIAT$ ratio increases, the average number of users in the system increases, and so a user has to travel farther in order to find an access point with a sufficiently low load. Therefore, the distance traveled increases with an increase in the $MIST/MIAT$ ratio.

The remaining graphs in Figure 4 indicate the stability of the system. Figure 4(d) shows the system Nash time. As seen in Figure 4(d), the system is always in Nash equilibrium when users exit simultaneously from the system; we proved this in Section IV-B. Under a dynamic exit model, this is no longer the case. We observe that the system Nash time decreases with an increase in the $MIST/MIAT$ ratio. To explain this behavior, we first analyze the effect of user arrivals and exits on system equilibrium. System equilibrium is never disturbed by the arrival of a user (note that we do not model diversity of individual workloads and α values in these simulations). The exit of a user reduces the load at the corresponding access point and can therefore disturb the equilibrium by motivating other users to move to that access point. The system equilibrium is subsequently restored either when a new user arrives and associates with the same access point, thereby increasing its load again, or when more exits occur at other access points and proportionately reduce their loads.

Consider that the exit of user q from access point j motivates user p to move from access point k to access point j . Assuming that user p adheres to her original selection, equilibrium is restored when either a new user arrives at access point j , or when an existing user from access point k exits the system. Since access point j is preferred by user p , it is highly likely that a new user also prefers the same access point. Therefore, the chance that equilibrium is restored within the next one or two user arrivals is high. On the other hand, from among the various access points in the system, the probability that the next exit is from access point k is low. Therefore, equilibrium is not likely to be restored soon if the exit of user q is followed by more user exits and no arrivals. Moreover, the subsequent exits may create further imbalance in the load distribution, thereby making it harder for equilibrium to be restored. In summary, the system equilibrium is likely to be restored quickly when a user exit is soon followed by a user arrival, and equilibrium remains disturbed for a longer duration when multiple exits occur in succession.

When the value of the $MIST/MIAT$ ratio is low, user arrivals and exits are fairly evenly intermingled throughout the duration of the simulation. In other words, a user exit is likely to be soon followed by a user arrival. As a result, equilibrium is restored quickly following any disruptions and the system Nash time is high. As $MIST/MIAT$ increases, user arrivals tend to be clustered in the earlier part of the simulation, while exits are clustered in the later part. Therefore, the exit of a user is less likely to be followed by an arrival, and system equilibrium disrupted by the exit does not get restored quickly, thereby reducing the system Nash time.

Next, we examine the trend in system Nash time with increasing α . At $\alpha = 0$, the system is always in a state of Nash equilibrium. An initial increase in the value of α causes system

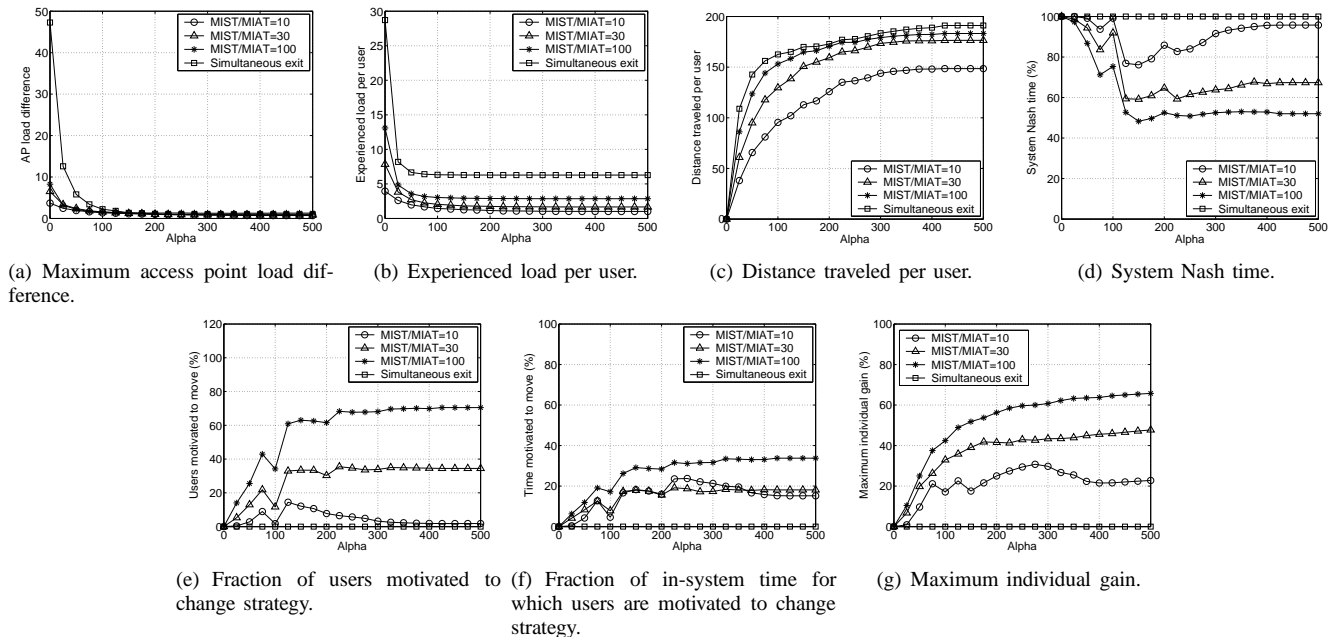


Fig. 4. Performance results with dynamic arrivals and exits.

Nash time to decrease since users attach greater importance to access point load and therefore have more motivation to move. At $\alpha = 100$, the system Nash time increases abruptly. Since this α value equals the inter-access-point distance (100m), users can select from a wider range of access points when they initially enter the system. As a result, the initial assignments are more spread out and therefore more stable. Similar behavior is observed at $\alpha = 200$. When $MIST/MIAT = 10$, a continued increase in α causes the stability of the system to increase significantly. Since on an average there are fewer users in the system, by initially moving longer distances, users distance themselves from each other and associate with lightly-loaded access points, thereby reducing the motivation to move later. Beyond a certain threshold, the increasing α value has no further impact on user decisions, and system Nash time remains constant. The higher the value of the $MIST/MIAT$ ratio, the greater the average number of users in the system, and therefore the smaller the threshold value of α beyond which system stability is unaffected.

Figure 4(e) indicates the average number of users that are motivated to change strategy, while Figure 4(f) indicates the average fraction of in-system time for which users are thus motivated. Both graphs show trends that are complementary to the system Nash time in Figure 4(d) and can be explained in a similar manner. As seen in Figure 4(f), users are motivated to change strategy for only up to 35% of their in-system time in the worst case, i.e. the access point selected by the myopic algorithm when the user enters the system remains the optimal access point for 65% or more of the user's in-system time. This is an important result, clearly indicating the stability and effectiveness of the access point selection method.

Finally, Figure 4(g) shows the maximum instantaneous individual gain that a user can obtain through a change of strategy. This value increases with an increase in α since more

weight is assigned to the access point load parameter in the cost calculation. The value also increases with an increase in the $MIST/MIAT$ ratio due to the clustering of user exits, which lead to a greater instantaneous load difference between access points. The maximum value of the metric is around 65% in our test scenarios.

In summary, our results show that myopic access point selection based on the load-distance trade-off significantly improves load distribution. The magnitude of improvement increases with an increase in the average number of users present in the system. System stability is high in the steady state but decreases when a large number of user exits tend to occur in succession. The access point selected by the user on entering the system remains the optimal access point for that user for 65% or more of the user's in-system time in all our test scenarios. We also observe that the best trade-off between benefits, costs and stability is obtained at low values of α , i.e. between 30% to 50% of the inter-access-point distance. We have conducted more experiments to verify this behavior by increasing the inter-access-point distance to 150m and proportionately increasing the terrain size. The results of these experiments show very similar trends and validate our inference.

VII. CONCLUSION

Relocation of users from congested areas to less-crowded areas of a wireless network could be a simple and effective solution to balance network load distribution and thereby improve the distribution of bandwidth among users. However, the effectiveness and stability of such a system is questionable, due to the selfish nature of users and the lack of prior knowledge about future changes in user distribution, and must be analyzed before this solution can be deployed in real wireless networks.

In this paper, we presented an analysis of this solution using a game-theoretic approach. To our knowledge, this is the first such analysis of this idea. We modeled the system as a game and analyzed the resulting Nash equilibria. We found that, in a system consisting of homogeneous users and a simple arrival/exit model, a myopic distributed strategy leads this game to a Nash equilibrium in just a single iteration. When diversity is introduced among the users, the myopic algorithm is no longer guaranteed to produce a Nash equilibrium. However, the maximum individual gain that a user can obtain through a unilateral change of strategy is bounded, and this bound is proportional to the extent of diversity. In other words, the lower the diversity among the users, the greater the likelihood that the system will remain stable.

We used simulation to analyze the effectiveness and stability of myopic access point selection when users arrive and depart in a dynamic fashion. We found that the approach gives greater benefit when more users are present in the system on an average. The initial access point selection made by a user remains the optimal selection for 65% or more of the user's in-system time, clearly indicating the effectiveness of the approach. The system stability is high when user arrivals and exits are evenly intermingled; this is likely to be the case in several deployment scenarios, such as college campuses and shopping malls. The best trade-off between costs, benefits and stability is obtained when the value of α is between 30% and 50% of the inter-access-point distance. At these values of α , users can significantly improve experienced load by traveling relatively short distances.

Our work validates the effectiveness of users changing location on entering the system to improve received bandwidth. We provide a detailed analysis of the behavior of a system based on this idea in various usage scenarios. Such a system is simple to deploy and can bring improvement with relatively little effort from users. The magnitude of effectiveness and stability depends upon the characteristics of the deployment scenario. In future work, we plan to empirically examine how this idea performs under different arrival/exit patterns and when the user population is heterogenous. We would also like to explore more sophisticated non-linear modeling of user preferences, and other algorithms for access point selection under more complex system models.

VIII. ACKNOWLEDGEMENTS

This work was supported in part by NSF Career Award CNS-0347886 and by NSF NeTS Award CNS-0435527.

IX. BIOGRAPHIES

Kimaya Mittal (formerly Sanzgiri) is a Research Scientist working with the Communications Group at Citrix Online. She completed her Ph.D. in Computer Science at the University of California, Santa Barbara, in December 2006. The work published in this paper was completed as part of her Ph.D. research. Kimaya received her B.E.(Hons.) in Computer Science from the Birla Institute of Technology and Science (BITS), Pilani, India, in 1999. Her current research interests are in

the areas of networking, distributed systems, and network protocols.

Elizabeth M. Belding is an Associate Professor in the Department of Computer Science at the University of California, Santa Barbara. Elizabeth's research focuses on mobile networking, specifically mesh networks, multimedia, monitoring, and advanced service support. She is the founder of the Mobility Management and Networking (MOMENT) Laboratory (<http://moment.cs.ucsb.edu>) at UCSB. Elizabeth is the author of over 70 papers related to mobile networking and has served on over 50 program committees for networking conferences. Elizabeth is the recipient of an NSF CAREER award, and a 2002 Technology Review 100 award, awarded to the world's top young investigators. See <http://www.cs.ucsb.edu/ebelding> for further details.

Subhash Suri is a professor of Computer Science at the University of California, Santa Barbara. His current research interests are algorithms, networked sensing, data streams, computational geometry, and game theory. He has published more than 150 research papers, and has been granted several patents for his algorithmic innovations. His research has been supported by several grants from the National Science Foundation.

REFERENCES

- [1] A. Akella, S. Seshan, R. Karp, S. Shenker, and C. Papadimitriou. Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP. In *ACM Sigcomm*, Pittsburgh, PA, Aug 2002.
- [2] L. Anderegg and S. Eidenbenz. Ad-hoc VCG: a Truthful and Cost-Efficient Routing Protocol for Mobile Ad hoc Networks with Selfish Agents. In *ACM MobiCom*, San Diego, CA, Sep 2003.
- [3] A. Balachandran, P. Bahl, and G. Voelker. Hot-Spot Congestion Relief in Public-Area Wireless Networks. In *IEEE WMCSA*, Callicoon, NY, Jun 2002.
- [4] A. Balachandran, G. Voelker, P. Bahl, and P. V. Rangan. Characterizing User Behavior and Network Performance in a Public Wireless LAN. In *ACM Sigmetrics*, Marina Del Rey, CA, Jun 2002.
- [5] M. Balazinska and P. Castro. Characterizing Mobility and Network Usage in a Corporate Wireless Local-Area Network. In *MobiSys*, San Francisco, CA, May 2003.
- [6] B. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz. Characterizing Selfishly Constructed Overlay Routing Networks. In *IEEE Infocom*, Hong Kong, Mar 2004.
- [7] S. Eidenbenz, V. Kumar, and S. Zust. Equilibria in Topology Control Games for Ad hoc Networks. In *DialM-POMC*, San Diego, CA, Sep 2003.
- [8] A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and S. Shenker. On a Network Creation Game. In *ACM PODC*, Boston, MA, Jul 2003.
- [9] T. Henderson, D. Kotz, and I. Ayzov. The Changing Usage of a Mature Campus-wide Wireless Network. In *ACM MobiCom*, Philadelphia, PA, Sep 2004.
- [10] R. K. Lam, D. M. Chiu, and J. C. S. Lui. On the Access Pricing and Network Scaling Issues of Wireless Mesh Networks. In *IEEE ICDCS*, Lisboa, Portugal, July 2006.
- [11] R. K. Lam, D. M. Chiu, and J. C. S. Lui. On the Access Pricing and Network Scaling Issues of Wireless Mesh Networks. *IEEE Transactions on Computers*, 56(11), Nov 2007.
- [12] M. Mavronikolas and P. Spirakis. The Price of Selfish Routing. In *Symposium on the Theory of Computing*, Heraklion, Greece, Jul 2001.
- [13] X. Meng, S. Wong, Y. Yuan, and S. Lu. Characterizing Flows in Large Wireless Data Networks. In *ACM MobiCom*, Philadelphia, PA, Sep 2004.
- [14] J. Musacchio and J. Walrand. WiFi Access Point Pricing as a Dynamic Game. In *7th INFORMS Telecommunications Conference*, Boca Raton, FL, Mar 2004.
- [15] C. Papadimitriou. Algorithms, Games and the Internet. In *Symposium on the Theory of Computing*, Heraklion, Greece, Jul 2001.

- [16] L. Qiu, Y. Yang, Y. Zhang, and S. Shenker. On Selfish Routing in Internet-Like Environments. In *ACM Sigcomm*, Karlsruhe, Germany, Aug 2003.
- [17] S. M. Ross. *Introduction to Probability and Statistics for Engineers and Scientists*. Harcourt Academic Press, 2nd edition, 2000.
- [18] T. Roughgarden and E. Tardos. How Bad is Selfish Routing? In *FOCS*, Redondo Beach, CA, Nov 2000.
- [19] K. Sanzgiri and E. Belding-Royer. Leveraging Mobility to Improve Quality of Service in Mobile Networks. In *MobiQuitous*, Boston, MA, Aug 2004.
- [20] C. Saraydar, N. Mandayam, and D. Goodman. Efficient Power Control via Pricing in Wireless Data Networks. *IEEE Transactions on Communication*, 50(2):291–303, Feb 2002.
- [21] M. Satyanarayanan. Pervasive Computing: Visions and Challenges. *IEEE Personal Communications*, 8(4):10–17, Aug 2001.
- [22] V. Srinivasan, P. Nuggehalli, C. Chiasserini, and R. Rao. Co-operation in Wireless Ad Hoc Networks. In *IEEE Infocom*, San Francisco, CA, Mar 2003.
- [23] S. Suri, C. Toth, and Y. Zhou. Selfish Load Balancing and Atomic Congestion Games. In *ACM SPAA*, Barcelona, Spain, Jun 2004.
- [24] W. Wang, X. Li, and Y. Wang. Truthful Multicast Routing in Selfish Wireless Networks. In *ACM MobiCom*, Philadelphia, PA, Sep 2004.