

Optimizing Placement of Wireless Access Points

Balaji Rengarajan and Gustavo de Veciana
Dept. of Electrical and Computer Engineering
The University of Texas at Austin
Austin, Texas 78712
Email: (balaji, gustavo)@ece.utexas.edu
Telephone: (512) 471-2358

Abstract

Wireless LANs have become increasingly popular for providing high data rate network access. Ideally, a large-scale wireless LAN is designed such that there are no gaps in its coverage. It is also desirable to minimize the number of access points used in order to minimize the cost of equipment, installation and maintenance. In this paper, we propose a grid-based approximation algorithm to compute the placement of access points. The algorithm minimizes the number of access points required while ensuring that the received SNR at each location is sufficient to meet the offered load at that location. The algorithm can be used when the offered load is spatially non-homogenous or when access points can be operated in different modes. The proposed algorithm is guaranteed to find a solution with at most $O(c_{opt} \log c_{opt})$ access points, if there exists a solution to the access point placement problem using c_{opt} access points.

Index Terms

Wireless networks, network design, access point placement, network provisioning, load-dependent design

Optimizing Placement of Wireless Access Points

Balaji Rengarajan and Gustavo de Veciana
Dept. of Electrical and Computer Engineering
The University of Texas at Austin
Austin, Texas 78712
Email: (balaji, gustavo)@ece.utexas.edu

Abstract—Wireless LANs have become increasingly popular for providing high data rate network access. Ideally, a large-scale wireless LAN is designed such that there are no gaps in its coverage. It is also desirable to minimize the number of access points used in order to minimize the cost of equipment, installation and maintenance. In this paper, we propose a grid-based approximation algorithm to compute the placement of access points. The algorithm minimizes the number of access points required while ensuring that the received SNR at each location is sufficient to meet the offered load at that location. The algorithm can be used when the offered load is spatially non-homogenous or when access points can be operated in different modes. The proposed algorithm is guaranteed to find a solution with at most $O(c_{opt} \log c_{opt})$ access points, if there exists a solution to the access point placement problem using c_{opt} access points.

I. INTRODUCTION

Wireless LANs (WLANs) were originally intended to provide connectivity in locations where using conventional wired LANs is not feasible. However, in the last few years, WLANs have become increasingly popular to provide high data rate network access. WLANs are typically used in offices, campus environments, shopping malls, airports etc. WLANs generally operate in the unlicensed industrial, scientific and medical (ISM) bands at 2.4 GHz, and 5 GHz.

Ideally, a large-scale WLAN has to be designed such that there are no coverage gaps, i.e., the received signal strength at every location in the target space must exceed a certain threshold. This objective is achieved by suitably placing wireless access points. It is also desirable to minimize the number of access points used in order to minimize the cost of equipment, installation and maintenance. Also, overprovisioning the WLAN by using more access points than necessary would only increase the overlap between coverage areas of access points using the same radio channel possibly degrading performance.

Wireless LAN design is complicated because the RF characteristics of the particular locations have to be considered. To date, WLANs are usually deployed in an ad-hoc manner, usually based on an educated guess by the person installing the access points followed by manual site surveys to ensure adequate coverage. A description of a typical WLAN deployment process can be found in [1], [2]. In this paper, we present an approximation algorithm to automate the placement of access points so as to achieve coverage while minimizing the number of access points required. We assume a CAD tool is available to model and/or simulate the environment's propagation characteristics [3].

The rest of this paper is organized as follows. We briefly describe the relevant features of the WLAN standards considered in this paper in Section II. Section III provides a brief overview of some other proposed methods for coverage planning in WLANs. The problem setup and the assumptions made are explained in Section IV. We review some relevant definitions from computational geometry in Section V. In Section VI, we formulate the access point placement problem as a set cover problem. The algorithm itself is described in Section VII. A brief analysis of the complexity of the proposed algorithm is presented in Section VIII and we conclude in Section IX.

II. WIRELESS LAN STANDARDS

IEEE 802.11 is the most prominent standard among existing wireless LAN technologies. The IEEE 802.11b standard uses direct sequence spread spectrum to support data rates up to 11 Mbps in the 2.4 GHz band. The IEEE 802.11a standard specifies an OFDM physical layer that splits an information signal across 52 separate subcarriers to provide transmission of data at rates up to 54 Mbps in the 5 GHz band. Both standards allow only a discrete set of data rates to be used. An automatic rate fallback function selects the data rate to be used depending on the signal to noise ratio observed at the receiver. The IEEE 802.11b standard operates at 1, 2, 5.5 and 11 Mbps depending on the observed signal to noise ratio while the IEEE 802.11a standard provides rates of 6, 9, 12, 18, 24, 36, 48, and 54 Mbps.

III. BACKGROUND

The problem of determining 'optimal' access point locations is a hard one in general. Hence [4] presented an approximation algorithm based on simulated annealing. [5] and [6] propose two approximation algorithms based on genetic algorithms, while an algorithm based on the gradient descent technique is presented in [7]. The convergence properties of these algorithms depends strongly on the choice of the initial solution and the simulation parameters. Since a good initial solution and simulation parameters are not always immediately apparent, the algorithms often suffer from slow convergence and sometimes, the algorithms do not converge at all. So, no bound can be provided on the number of iterations required to reach the final solution. Also, no bound is provided on how far the final solution is from the optimal solution. In [8], a simple heuristic algorithm based on pruning and neighborhood search is proposed. The proposed algorithm, however, only finds a local optimum. To find a solution close to the global optimum,

this solution has to be used as the starting condition for one of the other algorithms above. So, this algorithm suffers from the same drawbacks as the others.

IV. THE ACCESS POINT PLACEMENT PROBLEM

When deploying a wireless network, ideally we would like to ensure that all relevant locations in the target area are covered. In a simple scenario, this could mean that the received SNR at any point in the target area must exceed a specified threshold. However, in reality, the offered load may not be uniform. Indeed, it may not be necessary to support a high data rate at all locations. For example, one might want to ensure that users in a conference room have access to high capacity, while it is sufficient to design for lower capacity in hallways and corridors. So, we would like to find the placement for the minimum number of access points that provide the 'required' capacity at each location in the target area. In general, the above problem is very difficult to analyze. In this paper, we consider the following discretized access point placement problem.

The target area is overlayed by a square grid, where the length of the side of each square is δ . Each square in the grid is indexed and we let X denote the index set of squares in the grid.

The wireless LAN designer specifies a minimum data rate for each square in the grid that the wireless LAN should be able to support. This is different from the coverage problem usually considered in the literature. Here, not only do we have to ensure that all areas are covered, but we also have to guarantee that the SNR at each point is high enough to meet the demand. For coverage to be achieved, the received signal strength at each square has to be high enough to support the specified data rate. Thus, we can convert the data rate requirements to a received signal strength requirement in each square.

The possible locations of the access points would also in general be constrained, as installation, maintenance and administration costs would make some locations infeasible. Generally, access points can be placed only in the corners of rooms, or on walls lining corridors etc. So, the LAN designer specifies a set of candidate locations for the placement of access points. Let n denote the number of candidate access point locations and A , the set of candidate access point locations.

For each $a \in A$, we let the set $R_a \subseteq X$ denote the region covered by an access point placed at a . A grid square, $x \in X$ is said to be covered by an access point at location a , i.e. $x \in R_a$, if the strength of the received signal from the access point satisfies a specified condition. For example, the condition might be that the received signal strength exceeds a specified threshold $s(x)$ at the midpoint of x or we might require that the mean received signal strength in the square x exceeds a specified threshold. Alternate conditions evaluated on the square might also be considered. Different methods to determine the coverage area of an access point are discussed below.

A. Determining the Coverage Areas

Predicting the received signal strength at a location given the location of the transmitter is a non-trivial problem. In free space, the intensity of radio waves attenuates as the inverse square of the distance travelled. However, in the presence of a large number of obstacles like walls, furniture, several other mechanisms such as reflection, refraction, diffraction and scattering [3] come into play. The received signal strength depends on the electrical characteristics and the dimensions of the obstacles.

There are a number of tools that have been developed based on various modelling techniques to predict the signal strength at a location for a particular placement of an access point. The most popular models used for signal strength prediction are empirical models and ray-tracing based models. These models are discussed in detail in [3].

For each of the candidate access point locations specified by the WLAN designer, the coverage area can be determined by using one of the above tools to compute the received signal strength at the center of each square and then determining if the access point covers the square by comparing this signal strength to the minimum required signal strength, for that square. Thus, given the SNR requirements at each grid square, we can define the coverage area R_a corresponding to an access point at location a . We assume that all access points transmit at the maximum transmit power, P_{\max} .

Figure 1 depicts an example scenario. Figure 1(a) shows an example SNR specification in a convention center, while Figure 1(b) provides an example of the coverage area for a particular access point location.

Thus, we define the discretized access point placement problem as follows:

Problem 1: Given X , the set of grid squares, A , the set of candidate access point locations and coverage areas R_a for each access point location in A ; find the smallest set $L \subseteq A$ that satisfies the following condition: $\forall x \in X, \exists l \in L$ such that x is covered by an access point placed at l .

V. RELEVANT DEFINITIONS

We first review some basic definitions and results from computational geometry. We will use, for the most part, the terminology introduced in [9]. A set system (X, \mathcal{R}) consists of a set X along with a collection \mathcal{R} of subsets of X . \mathcal{R} is called the range space.

A *set cover* is a subcollection $\mathcal{T} \subseteq \mathcal{R}$ such that $\bigcup_{T \in \mathcal{T}} T = X$. The size of the cover is the number of sets in \mathcal{T} , $|\mathcal{T}|$. The *set cover problem* is that of finding the set cover of minimum size.

A related problem is the *hitting set problem*. A *hitting set* is a subset $U \subseteq X$ such that U has a non-empty intersection with every set $R \in \mathcal{R}$, that is, $R \cap U \neq \emptyset, \forall R \in \mathcal{R}$. The size of the hitting set is given by the number of elements in U , $|U|$. The *hitting set problem* is that of finding the hitting set of minimum size.

In fact these two problems are duals of each other. Consider the dual set system $(\mathcal{R}, \mathcal{X}^*)$, where $\mathcal{X}^* = \{\mathcal{R}_x^* : x \in X\}$ and

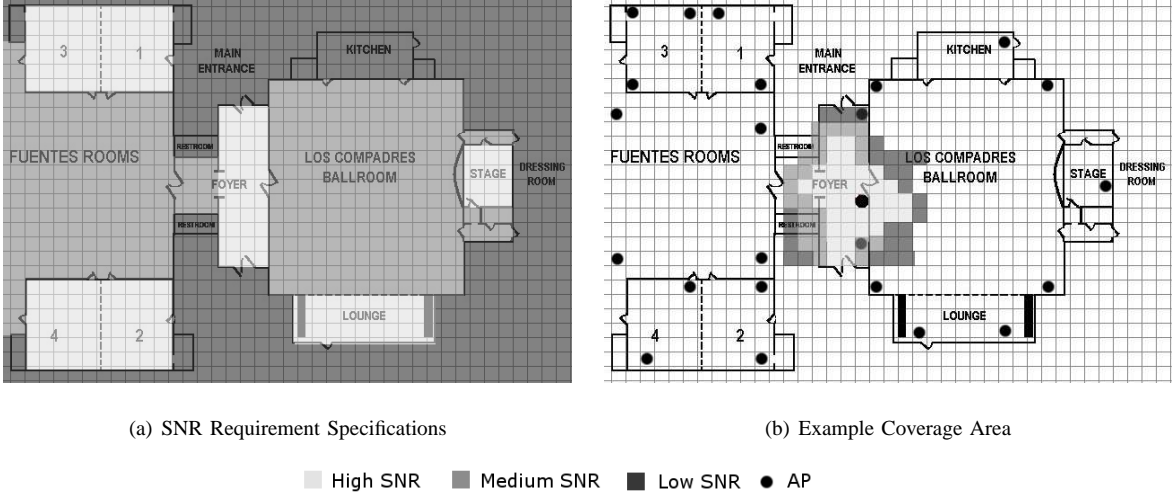


Fig. 1.

$\mathcal{R}_x^* = \{R : R \in \mathcal{R}, x \in R\}$. Thus, \mathcal{X}^* is a collection of collections of sets, each of which corresponds to a certain $x \in X$. The collection corresponding to x consists of all the sets in \mathcal{R} that hit x . The smallest hitting set for this dual set system and the smallest set cover for the original set system are identical.

Given $Y \subseteq X$, all the subsets of Y obtained as the intersection of Y and $R \in \mathcal{R}$ form a set system called the set system induced by \mathcal{R} on Y , denoted by $\mathcal{R}|_Y$. We say Y is shattered by \mathcal{R} if $|\mathcal{R}|_Y| = 2^Y$. (X, \mathcal{R}) is said to have VC-dimension d , if d is the smallest integer such that no $d + 1$ point subset $Y \subseteq X$ can be shattered [10].

If a subset $N \subseteq X$ has a non-empty intersection with each set $R \in \mathcal{R}$ of size greater than $\epsilon|X|$, it is called an ϵ -net [11]. This definition can be generalized by putting an additive weight function w on the power set of X . The additive weight function is defined as follows. A positive weight, $w(x)$ is associated with each $x \in X$. The weight of any subset of X is then defined to be the sum of the weights of the elements in the subset. Now, $N \subseteq X$ is said to be an ϵ -net if it has a non-empty intersection with all sets of \mathcal{R} whose weight is greater than $\epsilon w(X)$.

VI. FORMULATION AS A SET COVER PROBLEM

X denotes the index set of squares in the grid and A , the set of candidate access point locations. $R_a \subseteq X$ denotes the set corresponding to the coverage area of an access point placed at a . The discretized access point placement problem, Problem 1 can be formulated as a set covering problem.

Let \mathcal{R} be the collection containing the sets representing the coverage areas of the candidate access point locations, that is $\mathcal{R} = \{R_a : a \in A\}$. Consider the set system (X, \mathcal{R}) generated by the sets defined above. The minimum set cover consists of sets belonging to \mathcal{R} , each of which corresponds to a particular access point location. Thus the solution to the access point placement problem is the set of candidate

locations corresponding to the sets in the minimum set cover. The following problem is thus equivalent to Problem 1.

Problem 2: Given X , the set of grid squares and \mathcal{R} , the collection of coverage areas; find \mathcal{C} , the minimum set cover of the set system (X, \mathcal{R}) .

The dual set system is denoted by $(\mathcal{R}, \mathcal{X}^*)$. \mathcal{X}^* is a collection of collections of sets, $\mathcal{R}_x^*, x \in X$. \mathcal{R}_x^* is a collection of sets representing the coverage areas of candidate access point locations, that cover x . The solution to the minimum hitting set problem for the dual set system corresponds to the solution to the minimum set cover problem in the original set system. Notice that the minimum hitting set of the dual set system is a subcollection of \mathcal{R} . Thus, the following problem is the dual of Problem 2.

Problem 3: Given X , the set of grid squares and \mathcal{R} , the collection of coverage areas; find \mathcal{H}^* , the minimum hitting set of the dual set system $(\mathcal{R}, \mathcal{X}^*)$.

Table I summarizes the notation that is used in this paper.

X	Set of squares in the grid
x	a particular square in the grid
δ	Length of the side of a square in the grid
A	Set of candidate access point locations
n	Number of candidate access point locations
a	$a \in A$, a particular candidate location
R_a	Coverage region of access point at a
\mathcal{R}	$\{R_a : a \in A\}$, collection of coverage regions
$s(x)$	SNR threshold at square x
\mathcal{R}_x^*	Collection containing sets R_a , that cover x
\mathcal{X}^*	Collection containing collections $\mathcal{R}_x^*, \forall x \in X$
\mathcal{C}	Minimum set cover of (X, \mathcal{R})
\mathcal{H}^*	Minimum hitting set $(\mathcal{R}, \mathcal{X}^*)$
d	VC-dimension
c	Number of access points in a candidate set cover
c_{opt}	Number of access points in the minimum set cover

TABLE I
SUMMARY OF NOTATION

In the next section, we present an approximation algorithm based on an algorithm to find the minimum hitting set of a set system presented in [9]. The proposed algorithm tries to pick

the smallest number of access points that satisfy the coverage requirements. The following theorem characterizes the VC-dimension of the set system (X, \mathcal{R}) and will be necessary to prove the correctness of the proposed algorithm and to analyze its complexity.

Theorem 6.1: Given n candidate access point locations, the VC-dimension of the set system (X, \mathcal{R}) defined above is bounded by $\log_2(n)$ and the VC-dimension of the dual set system is bounded by $2n$.

Proof: Since the number of sets in \mathcal{R} equals n , the cardinality of $\mathcal{R}|_Y$ is at most n , thus if $|Y| > \log_2(n)$, Y cannot be shattered by \mathcal{R} . Also, the VC-dimension of the dual set system $(\mathcal{R}, \mathcal{X}^*)$ is less than 2^{d+1} , where d is the VC-dimension of the primal set system [12]. ■

VII. THE ALGORITHM

The algorithm is based on using an ϵ -net finder and a verifier. The net-finder and the verifier are discussed in detail in the sequel. The algorithm starts with a very low estimate of the number of access points in the minimum set cover, and at each stage tries to determine whether a set cover of that size exists. The ϵ -net finder is used at each stage to select candidate set covers. The verifier is used to determine if the net finder has found a set cover. If the algorithm determines that no set cover can be constructed using the estimated number of access points, the estimate for the required number of access points is increased and the algorithm continues until a set cover is found.

The weights of the collections $\mathcal{R}_x \in \mathcal{X}^*$ are computed as follows. Weights are assigned to each of the sets $R_a \in \mathcal{R}$. The weight of a collection, \mathcal{R}_x is then computed as the sum of the weights of all the sets in that collection.

Assume that we have a guess c for the minimum number of access points required. The algorithm initializes the value of c to 1 and the weights of all sets in \mathcal{R} to 1. The algorithm then tries to find $O(c \log c)$ access point locations that would satisfy the design requirements, i.e., it tries to find a hitting set of $(\mathcal{R}, \mathcal{X}^*)$ with $O(c \log c)$ elements.

At each iteration, the algorithm finds a weighted $\frac{1}{2c}$ -net of $(\mathcal{R}, \mathcal{X}^*)$. Thus, any member of \mathcal{X}^* that is not hit can have weight which is at most $\frac{1}{2c}$ of the total weight of \mathcal{X}^* . If the $\frac{1}{2c}$ -net found in the previous step does not correspond to a hitting set, then there exists a collection $\mathcal{R}_{x_w} \in \mathcal{X}^*$ that has no sets in common with the $\frac{1}{2c}$ -net, i.e. its intersection with the $\frac{1}{2c}$ -net is empty. Such a collection is called a witness and corresponds to a grid square x_w that is not covered if access points are placed at the locations corresponding to the coverage areas in the $\frac{1}{2c}$ -net. If the sets in the $\frac{1}{2c}$ -net do not cover X , then the verifier is used to pick one such witness and the weights of all the members of the witness are doubled.

It is shown in [9], that, if there is a hitting set of $(\mathcal{R}, \mathcal{X}^*)$ of size c , the doubling process cannot iterate more than $4c \log_2(\frac{n}{c})$ times. Thus, if after $4c \log_2(\frac{n}{c})$ iterations, the algorithm has not found a hitting set, the algorithm determines that a hitting set of size c does not exist. Then, it increments

its guess for c , the number of access points required, resets the weights of the sets in \mathcal{R} and tries to find a hitting set.

This process is repeated until a hitting set is found. The value of c can, say, be doubled at the end of each pass to ensure that the algorithm speedily finds a solution. The algorithm is presented in Table II.

A. The ϵ -net Finder

Lemma 7.1: The set system $(\mathcal{R}, \mathcal{X}^*)$ admits a $\frac{1}{r}$ -nets of size $O(n \log_2(nr))$.

Proof: It is proved in [11], that set systems of VC-dimension d admit a $\frac{1}{r}$ -net of size $O(d \log_2(dr))$, and Theorem 6.1 provides a bound on the VC-dimension of the set system under consideration. ■

Algorithms to find ϵ -nets for set systems of finite VC-dimension are described in [11], [13] and [14]. It is proved in [11] that, for a set system of VC-dimension d , choosing a random sample of size $O(d \log c)$ according to the distribution induced by the weights, results in a $\frac{1}{c}$ -net with high probability. So, in step (i), we can pick a $\frac{1}{2c}$ -net of $(\mathcal{R}, \mathcal{X}^*)$ by choosing at random $O(nc \log c)$ sets from \mathcal{R} according to the distribution induced by the weights. If this sample does not yield a $\frac{1}{2c}$ -net, i.e., if we find in step (iii), a square x_w that is not covered such that the weight of the corresponding set in \mathcal{X}^* is greater than a $\frac{1}{2c}$ fraction of the total weight of \mathcal{X}^* , then we discard the random sample and we pick another random sample and repeat the steps.

B. The Verifier

In steps (ii) and (iii), we need to check if the ϵ -net found in step (i) is a hitting set for $(\mathcal{R}, \mathcal{X}^*)$. If it is a hitting set, then this corresponds to the set cover in (X, \mathcal{R}) and we have found a set of locations that meet the design requirements. Otherwise, we need to find a witness, i.e., an element of \mathcal{X}^* that is not hit, corresponding to a square on the grid that is not covered.

As we find the coverage areas of the candidate locations, we also build the following table. The table stores an n bit coverage-vector for each square in the grid. The i^{th} bit in the coverage-vector corresponding to square x_j is set to 1 if x_j is covered by R_i . The random sample is also represented by a n bit sample-vector, and the i^{th} bit in this vector is set to one if R_i is chosen. So, we can determine if a particular square is covered by computing the logical and of the coverage-vector and the sample vector and comparing the result to the zero-vector. Thus, the verifier checks to see if each square is covered and if the random sample is not a hitting set, returns the first square that it finds that is not covered.

VIII. ANALYSIS OF THE ALGORITHM

If there exists a solution to the access point placement problem, using c_{opt} access points, the algorithm is guaranteed to find a solution with at most $O(c_{\text{opt}} \log c_{\text{opt}})$ access points. Even if we double the guess for c at the end of each pass, the hitting set returned will be of size $O(4c_{\text{opt}} \log 4c_{\text{opt}})$.

- 1) Set $c = 1$
- 2) Do while $c \leq n$:
 - a) Assign weight 1 to all the sets $R \in \mathcal{R}$ or equivalently to each of the candidate access point locations.
 - b) For $k := 1$ to $4c \log_2 \left(\frac{n}{c}\right)$ do:
 - i) Using the ϵ -net finder, select a $\frac{1}{2c}$ -net of $(\mathcal{R}, \mathcal{X}^*)$
 - ii) Using the verifier, check to see if the chosen sets of \mathcal{R} cover X . If so, return the access point locations corresponding to the chosen sets of \mathcal{R} and exit.
 - iii) If not, find a witness square $x_w \in X$ that is not covered by any of the chosen sets.
 - iv) Double the weight of all the sets in \mathcal{R} (equivalent to doubling the weight of the corresponding location) that cover x_w
 - c) Increase c
- 3) Failure: No solution is possible.

TABLE II
THE ALGORITHM

We require the received signal strength to be calculated at the midpoint of each square in the grid, for each candidate location. If a propagation modelling tool using a ray-tracing based approach were to be used, this would be computationally demanding. But this computation is performed just once to generate the input to the algorithm. The time complexity of the ϵ -net finder, when finding a $\frac{1}{2c}$ -net is $O(c \log c)$. The time complexity of the verifier is $O(Nn)$, where N is the total number of squares in the grid. The time complexity of the algorithm itself when we double the value of c at the end of each pass is $O\left(c_{opt} \log \left(\frac{n}{c_{opt}}\right) ((c_{opt} \log c_{opt}) + Nn)\right)$, where n denotes the number of candidate access point locations.

IX. FUTURE WORK AND CONCLUSION

We formulated the access point placement problem as a set cover problem. We proposed an algorithm based on an approximation algorithm for the set cover and hitting set problems that met the received signal strength constraints at all points in the area of interest while attempting to minimize the number of access points that need to be deployed.

The proposed algorithm does not take into account the effect of co-channel interference. The IEEE 802.11a and 802.11b standards allow for multiple frequencies of operation to be used. The co-channel interference can be minimized by appropriately selecting the frequency of operation of each access point. Also, the co-channel interference can be accounted for by increasing the required signal to noise ratio at each square by a small margin.

Also, the proposed algorithm can be used in the case where the access points can be operated in different modes that correspond to different coverage areas. The only modification necessary is to add multiple instances of each access point location to A corresponding to the different modes. The algorithm can then be used not only to select the optimal positions of the access points, but also to select the mode in which each access point should be operated. This would not preclude the case where the optimal solution has multiple instances of the same access point, but this would be unlikely in a practical situation as there will be a large amount of overlap between the coverage areas of an access point operating in different

modes, and the minimum set cover is unlikely to contain two sets with a large degree of overlap.

In the future, we intend to consider the scenario where there is a different cost associated with each candidate location. The cost could represent wiring and administration costs. The different costs could also be used to model access points that can operate in multiple modes, with different modes having different associated costs. We are also currently working to develop an algorithm that takes into consideration the effects of co-channel interference.

REFERENCES

- [1] A. Raniwala and T. Chiueh, "Deployment issues in enterprise wireless LANs," 2003.
- [2] A. Hills, "Large-scale wireless LAN design," *IEEE Commun. Mag.*, pp. 98–104, November 2001.
- [3] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall, 2002.
- [4] H. Anderson and J. McGeehan, "Optimizing microcell base station locations using simulated annealing techniques," in *Proceedings of IEEE Vehicular Technology Conference*, vol. 2, Stockholm, Sweden, June 1994, pp. 858–862.
- [5] K. Lieska, E. Laitinen, and J. L  htenm  ki, "Radio coverage optimization with genetic algorithms," in *Proceedings of PIMRC*, vol. 1, Boston, MA, September 1998, pp. 318–322.
- [6] L. Nagy, L. Farkas, and B. Hegyi, "Optimizing indoor base station locations," in *ECWT*, 2000.
- [7] H. D. Sherali, C. M. Pendyala, and T. S. Rappaport, "Optimal location of transmitters for micro-cellular radio communication system design," *IEEE J. Select. Areas Commun.*, vol. 14, no. 4, pp. 662–673, May 1996.
- [8] M. Kamenetsky and M. Unbehaun, "Coverage planning for outdoor wireless LAN systems," in *International Zurich Seminar on Broadband Communications*, 2002.
- [9] H. Br  nnimann and M. T. Goodrich, "Almost optimal set covers in finite vc-dimension," *Discrete and Computational Geometry*, vol. 14, no. 4, pp. 463–479, 1995.
- [10] V. N. Vapnik and A. Y.   ervonenkis, "On the uniform convergence of relative frequencies of events to their probabilities," *Theory Probab. Appl.*, vol. 16, pp. 264–280, 1971.
- [11] D. Haussler and E. Welzl, " ϵ -nets and simplex range queries," *Discrete Comput. Geom.*, vol. 2, pp. 127–151, 1987.
- [12] P. Assouad, "Densit   et dimension," *Ann. Institut Fourier*, pp. 232–282, 1983.
- [13] H. Br  nnimann, B. Chazelle, and J. Matou  ek, "Product range spaces, sensitive sampling, and derandomization," in *Proc. 34th Annu. IEEE Sympos. Found. Comput. Sci.*, 1993, pp. 400–409.
- [14] B. Chazelle and J. Friedman, "A deterministic view of random sampling and its use in geometry," *Combinatorica*, vol. 10, no. 3, pp. 229–249, 1990.