Site Specific Knowledge for Improving Frequency Allocations in Wireless LAN and Cellular Networks

Jeremy K. Chen[†], Theodore S. Rappaport[‡], and Gustavo de Veciana[‡]

[†]Qualcomm Flarion Technologies, Inc., Bridgewater, New Jersey, USA

[‡]Wireless Networking and Communications Group (WNCG), The University of Texas at Austin, USA

Emails: {jchen, wireless, gustavo}@ece.utexas.edu

Abstract—This paper is the first analytical work to exhibit the substantial gains resulting from applying *site specific knowledge* to frequency allocation in wireless networks. Two new site-specific knowledge-based frequency allocation algorithms are shown to outperform all other published work. Site specific knowledge refers to knowledge of building layouts, the locations and electrical properties of APs, users, and physical objects. We assume that a central network controller communicates with all APs, and has site specific knowledge which enables the controller to predict, *a priori*, the received power from any transmitter to any receiver. Optimal frequency assignments are based on predicted powers to minimize interference and maximize throughput. Our algorithms consistently yield high throughput gains irrespective of network topology, AP activity level, and the number of APs, rogue interferers, and available channels. Our algorithms outperform the best published algorithm by up to 3.68%, 8.95%, 13.6%, 15.1%, 25.8%, and 84.9% for 50, 25, 20, 15, 10, and 5 percentiles of user throughputs, respectively.

I. INTRODUCTION

Radio propagation characteristics are fundamentally site specific, since radio propagation mechanisms (e.g. penetration, reflection, and diffraction) are directly related to the locations, sizes, and electrical properties of physical objects in the surroundings. Site-specific channel prediction algorithms [1]–[3] use a building layout or a satellite map and compute path losses between any two locations in indoor or outdoor environments. The complexity of these prediction tools has been reduced, and computing power has increased, so that they can be implemented for real-time network management applications. Today, site specific knowledge is becoming much more available from blueprints, AutoCAD, and Google[TM] Maps, for example. In this paper we use site specific knowledge to improve on-going frequency allocation in WLANs formed by APs and clients. In WLANs, a number of orthogonal frequency channels are allocated, and each AP is allocated one channel. When the number of channels is limited relative to the number of APs, some APs inevitably use the same channel and induce co-channel interference. The same problem exists in cellular networks. Judicious channel reuse mechanisms are necessary to reduce interference, particularly for the case of mobile users, such as in enterprise voice over IP networks or in cellular networks.

A number of WLAN frequency allocation schemes have been proposed thus far. The work in [4] assumes each AP

has a different fixed traffic load, and defines the effective channel utilization of an AP as the fraction of time the channel is used for data transmission or is sensed busy due to interference from other APs; then, the maximum effective channel utilization among all APs is minimized. AP placement and frequency allocation are jointly optimized in [5] with the same objective of minimizing the max channel utilization as in [4]. The frequency allocation problem is modeled as a minimum-sum-weight vertex-coloring problem in [6] where vertices are APs, and the weight of each edge between two APs denotes the number of clients that are associated with either one of these two APs and are interfered by the other AP. The work in [7] minimizes the number of clients whose transmissions suffer channel conflicts; a client associated with an AP suffers conflicts if other clients or other APs interfere with the client or the AP under consideration. The definition of channel conflict in [7] is more comprehensive than those in [4]–[6]; the work in [7] has been shown to outperform [4]–[6].

However, none of [4]–[7] presents mechanisms to detect and reduce the negative impact from *rogue interferers*, which refer to intentional or unintentional RF interferers in noncooperative networks, microwave ovens, or other RF devices that also operate on the unlicensed bands as WLAN. Only the work in [8], [9] handle rogue interferers. In [8], each AP senses interference and independently selects a channel whose measured interference power is below a predefined threshold, without coordinating with other APs. The work in [9] assumes that clients periodically report in situ interference measurements to their associated APs, and presents three iterative algorithms that use the reported measurements to minimize interference seen by clients. In each iteration, these three different algorithms reduce the overall interference (computed by some weighted function defined in [9]) seen in a single cell, a group of nearby cells, or all cells, respectively, where a *cell* means an AP and its associated users. The second algorithm with local coordination among nearby APs has been shown to be the best among the three in [9].

Most traffic in WLANs is downlink [3]; hence, maximizing signal-to-interference-and-noise ratio (SINR) and downlink throughput as seen by *users* are key to proper network design. The work in [4], [5], [8] minimizes the interference at *APs* rather than minimizes that at *users*, as is done in [6], [7], [9]. The work in [6]–[8] use a binary model for interference, which has been shown to be inferior than the physical interference

model used in [9]. Since [9] also has the ability to deal with rogue interferers, [9] has been shown to outperform [4]–[8]. Despite their success, the measurement-based algorithms in [9] can still be improved if we assume that a central network controller has and uses site specific knowledge to optimize frequency allocation of each AP and each user. The advantage of using site specific knowledge is to predict *a priori* path loss between any pair of AP and user, when the user's location is obtained via GPS (Global Positioning System) or other known position location technologies[2]. The predicted path losses can help formulate a global optimization problem, thereby maximizing throughputs and saving power, etc. In this work we assume that the actual path loss between any transmitter and receiver can be perfectly predicted by the site specific knowledge. Though the environment affects the path losses, empirical results show that by modeling large fixed partitions and items in the environment (such as walls, book shelves, and cubicles), the *predicted* and the *measured* path losses show high agreement (zero mean error and standard deviation of 3-4 dB) [1]–[3]. The study of the effect of imperfect predictions of channel gains is an ongoing and future work.

Distributed measurement-based algorithms with the knowledge of APs' transmit powers via message exchanges can learn over time the path loss or received power between every transmitter and every receiver; examples of such algorithms are described in [12]. Nevertheless, the time needed for learning may be long when the number of interfering APs is large. Saving the learning time for measurement-based algorithms is a topic for ongoing and future work; in this paper, we choose to use site specific knowledge to predict *a priori* the interference power between any transmitter and any receiver.

Note that the central controller must know the active transmitters at any point in time in order to predict correct interference at all times; this information may be too costly to obtain, but time sampling may be done. Since downlink volume presently dominates WLAN traffic, this paper considers a case where all APs are actively transmitting downlink traffic. It is reasonable to assume that frequency allocation is optimized with respect to this most active case, since in this case, frequency allocation is most crucial for interference mitigation. Simulations in Section IV show that our algorithms also perform well in scenarios with both downlink and uplink traffic and with different levels of AP activity.

We present system models, notation, and assumptions, followed by the details of the two algorithms based on site specific knowledge. Then we show by simulations that our algorithms substantially outperform the others in [4]–[9].

II. SYSTEM MODEL, NOTATION, AND ASSUMPTIONS

Suppose M APs, indexed by $\mathbb{M} = \{1, 2, \ldots, M\}$, operate on K orthogonal frequency channels, indexed by

---

[2]Several indoor position location approaches, based on signal strength sensing, are widely known today and used in some WLANs [10], [11]. Other triangulation methods can also be used to locate a client. Modern cellular handsets are equipped with GPS chips or other position location technologies. State-of-the-art GPS can work not only outdoors but also indoors; various vendors, e.g. Metris and SnapTrack, provide indoor GPS solutions.

$\mathbb{K} = \{1, 2, \ldots, K\}$. We index users (or clients) by $\mathbb{L} = \{1, 2, \ldots, L\}$. We denote the identity of an AP and a client by a_m ($m \in \mathbb{M}$) and c_l ($l \in \mathbb{L}$), respectively. We assume for this work the locations of the APs and the clients do not vary with time, and no APs or users are at the same location, although the algorithms given here also apply for mobile APs and/or clients. We assume every user is associated with a single AP. Let f_m ($f_m \in \mathbb{K}$) denote the channel that a_m operates on, and let $\vec{f} = \{f_1, f_2, \ldots, f_M\}$ denote the channels of all M APs.

We assume that the central network controller periodically (say every 5 minutes) requires the APs to stop transmitting for a short duration of time (say, one second). In this duration, APs take turns in requiring all users associated with them to perform measurements of background interference, which refers to both the noise floor and rogue interference from RF devices outside the controlled network. Note that each user needs to measure the background interference for all available frequency channels. The users then feedback to APs these measured background interference. Site specific knowledge along with measurements of background interference make the estimations of SINR at users or APs much more accurate.

III. SITE-SPECIFIC KNOWLEDGE-BASED ALGORITHMS

A. The **Site-Specific SINR (SS-S)** Formulation

We shall consider optimizing a sum of utility functions for all the users' SINR, assuming all APs are actively transmitting downlink traffic (but not uplink traffic). That is, we optimize the following problem over $\vec{f} \in \mathbb{K}^M$, which is denoted *Site Specific SINR* or *SS-S* in the rest of this paper:

$$\max_{\vec{f} \in \mathbb{K}^M} \left\{ \sum_{l \in \mathbb{L}} U(\gamma_l) \,|\, \gamma_l = \frac{S_{c_l, a_m}}{P^i_{c_l} + \sum_{n: f_n = f_m, n \neq m} S_{c_l, a_n}}, \forall l \in \mathbb{L} \right\} \quad (1)$$

where a_m in (1) denotes the AP with which c_l is associated, $P^i_{c_l}$ denotes background interference power that c_l measures (as described in Section II), γ_l denotes the SINR at user c_l as shown in (1), S_{c_l, a_m} denotes the average received signal power from a_m to c_l. Note that the objective in (1) is not optimizing 'sum SINR', since such an objective may favor users that are closer to APs and may cause users which are further away to suffer low SINR. A fair SINR distribution can be achieved if we optimize the sum of utility functions in (1), where the utility function $U(\cdot)$ in (1) can be any function that is concave, continuously differentiable, and strictly increasing. For example, Mo and Walrand have proposed a class of utility functions that capture different degrees of fairness parameterized by q [13]:

$$U(\gamma_l) = \begin{cases} (1-q)^{-1} \gamma_l^{(1-q)}, & \text{if } q \neq 1 \\ \log \gamma_l, & \text{if } q = 1 \end{cases}, \gamma_l \in (0, \infty). \quad (2)$$

This family of utility functions is concave, continuously differentiable, and strictly increasing. Intuitively, as q increases, the degree of fairness increases, but the sum SINR decreases. A trade-off between sum SINR and fairness of the individual SINR of users that are further away from a serving AP can be

observed. By increasing the degree of fairness, we imply that users that are further from APs have higher SINR (which is needed to provide high throughput to distant users). The work in [13] shows that if $q \to \infty$, the formulation in (1) becomes a special case that achieves max-min fairness.

At max-min fairness, the degree of fairness is the highest; however, the sum SINR is the lowest. Simulation results in Section IV show that $q = 2$ may be a good parameter to capture this trade-off, but this remains a topic for further research.

As described in Section II, we assume that the APs and/or the users in the controlled network periodically measure the background interference; hence, $P_{c_l}^i$ in (1) is known. We assume that the central network controller has site specific knowledge and the locations of all APs and users, can predict signal power for any pair of AP and user, and can compute S_{c_l,a_n} for all c_l, a_n in the denominator of (1). Thus all the quantities in the optimization problem in (1) are known, yet measurement-based algorithms (such as the ones presented in [9]) do not know each individual component of S_{c_l,a_n} in (1) and thus are not able to solve (1) directly. Because the optimization in (1) is a combinatorial problem, there is no fast algorithm (polynomial-time) that can solve (1). Therefore, we propose an efficient heuristic in Section III-C that can find the locally optimal solution of (1); simulations show that the algorithm in III-C outperforms the measurement-based algorithms in [9], as well as all other frequency allocation algorithms in [4]–[8].

B. The **Site-Specific Rate (SS-R)** Formulation

The formulation in (1) in Section III-A strives to provision fair SINR across users. From the users' perspective, however, *throughput* may be a better metric than SINR for users' performance. Below we formulate another problem that aims at provisioning fair throughput across users, and this formulation may be denoted *Site Specific Rate* or *SS-R*.

$$\max_{\vec{f} \in \mathbb{K}^M} \left\{ \sum_{l \in \mathbb{L}} U(\chi_l) \mid \chi_l = \frac{r_l(\gamma_l)}{L_m}, \right. \tag{3}$$

$$\left. \gamma_l = \frac{S_{c_l,a_m}}{P_{c_l}^i + \sum_{n:f_n=f_m, n \neq m} S_{c_l,a_n}}, \forall l \in \mathbb{L} \right\} \tag{4}$$

where L_m denotes the number of clients that are associated with a_m, χ_l denotes the throughput of c_l from a_m (c_l is associated with a_m), $r_l(\gamma_l)$ denotes the long-term average data rate that c_l receives from a_m if c_l is the only user associated with a_m; r_l depends on the SINR seen at user c_l, i.e., γ_l. $r_l(\gamma_l)$ may also be viewed as the achievable capacity between c_l and a_m. We assume that the AP a_m evenly distributes its resource (e.g. time) amongst its L_m users and therefore has the denominator in (4). There are several ways to model $r_l(\gamma_l)$; for example, we may use Shannon capacity

$$r_l(\gamma_l) = \log_2 (1 + \gamma_l) \tag{5}$$

or an empirical model, e.g., such as introduced in [3], [14] to relate throughput to SINR:

$$r_l(\gamma_l) = T_{\max} \left(1 - e^{-A_e(\gamma_l - \gamma_0)} \right), \tag{6}$$

where the three constants $T_{\max}$, A_e, and γ_0 denote peak throughput, slope of throughput variation, and the cutoff SINR, respectively, as described in [3]. The model in (6) captures the downlink throughput of a client c_l when all other clients associated with the same AP are idle. In our simulation, we use a time division multiplexing model. Hence, at any point of time, an AP is sending data to only one client. The model in (6) is valid, as long as we multiply the throughput in (6) by the time fraction that a_m allocates to Client c_l.

C. A Local Optimization Algorithm for SS-S and SS-R

The optimization problems in (1) and (3) are combinatorial; solving them exhaustively requires exponential computation time (exponential in the number of APs). Hence, we present an iterative local optimization procedure that yields rapid and nearly-optimal solutions of (1); the same procedure can also solve (3). At the beginning of each iteration, a frequency allocation $\vec{f}$ is given, and at the end of the iteration, we find a better frequency allocation $\vec{g}$ that improves the objective in (1); $\vec{f}$ and $\vec{g}$ may differ in several elements, which means that the channels of several APs may change. During each iteration, we do the following steps. First, we select an AP, say a_m. We find $V-1$ other APs that produce the strongest interference on a_m, assuming these $V-1$ other APs and a_m are on the same channel; for example, $V = 7$ implies that we find 6 other APs that are in the vicinity of a_m so that they will likely interfere with a_m's clients the most. We try all possible K^V permutations of channels for these V APs, while fixing the channels at the other $M - V$ APs. We can find the best out of the K^V permutations so that (1) is maximized and is strictly larger than the value before this iteration; then we change the corresponding V elements in $\vec{f}$ and thus form $\vec{g}$. If these V APs have operated on the best channel allocation before this iteration, we have $\vec{f} = \vec{g}$; in this case, another AP (instead of a_m) and its $V-1$ neighboring APs will be selected to restart this iteration. This iterative algorithm runs until every set of V neighboring APs reaches the best frequency allocation. This iterative algorithm converges in a finite number of steps, since the number of channel permutations is finite, and each iteration strictly increases the objective in (1). In practice, one may limit the number of iterations or specify a minimum difference of weighted interference so that the iterations can be finished in a reasonable amount of time (say 1, 5, or 30 seconds). We expect that the channel allocation found by this local optimization algorithm will be close to the optimum if V is large enough, since the exhaustive search can explore more possible allocations with a larger V. The simulations in Section IV shows that this local optimization algorithm with $V = 7$ outperforms all other algorithms [4]–[8].

The algorithm proposed above solves the *SS-S* formulation in (1) and the *SS-R* in (3). When it is applied to solve *SS-S*, we refer to the algorithm as the *SS-S* algorithm; similarly, when

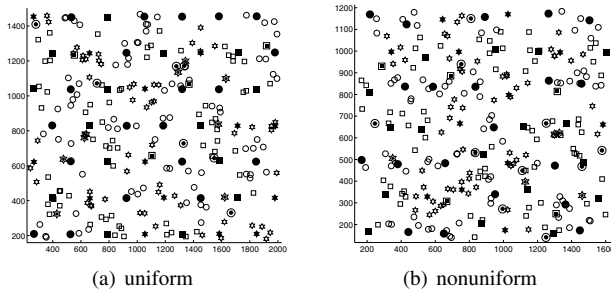(a) uniform   (b) nonuniform

Fig. 1. Frequency allocation examples for 49 APs on a 7-by-7 *nonuniform* or *uniform* topology. Three kinds of objects (squares, stars, and circles) signify three orthogonal frequency channels. Filled back objects denote 49 APs; hollow objects denote 196 users; double-layered objects with inner part filled with black denote 20 rogues. The units of X and Y axes are meters.
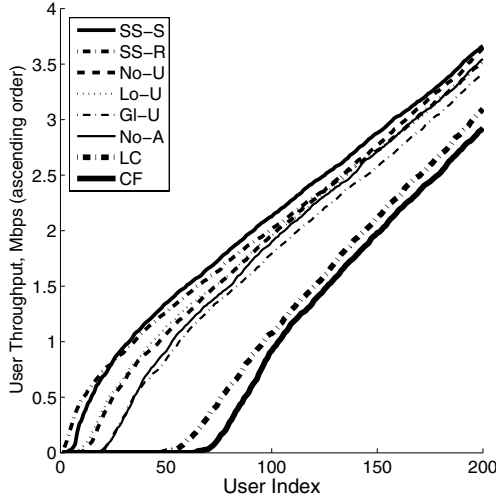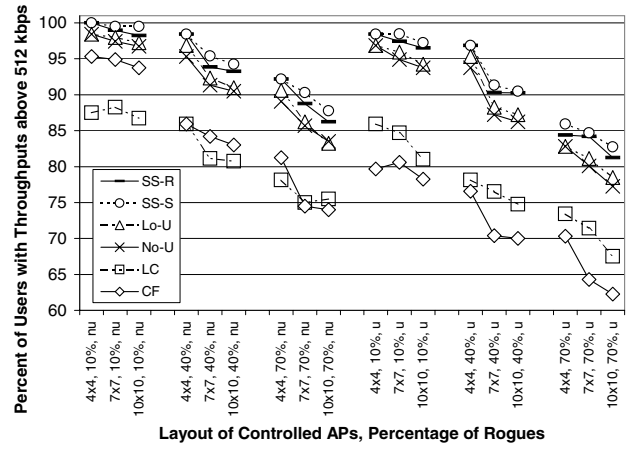

Fig. 3. Percent of users that have throughputs higher than 512 kbps. The x-axis represents the layout of controlled APs and the percentage of rogue APs compared to the controlled APs. *Nonuniform* and *uniform* AP layouts are denoted 'nu' and 'u', respectively.
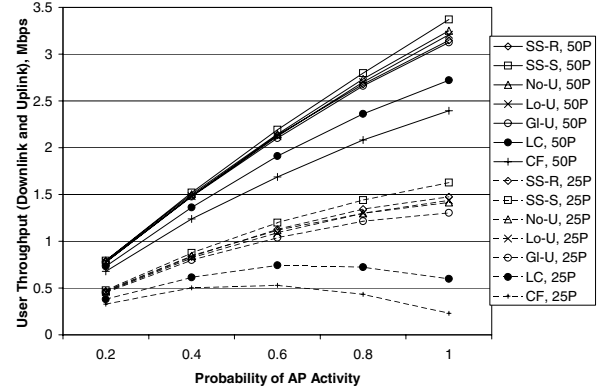

Fig. 2. *User throughput (in Mbps)* comparison with different levels of noise floor in a setting with APs on a uniform 10-by-10 layout, 400 users, and 10 rogue RF interferers. Only the 200 users with the lowest throughputs are shown.

Fig. 4. 50P and 25P signify 50 and 25 percentiles of users' throughputs respectively, including both downlink and uplink traffic.

the algorithm is used to maximize throughput (rate), we refer to it as the *SS-R* algorithm.

IV. SIMULATION SETUP AND RESULTS

The algorithm in [7], denoted as *CF*, has been shown to outperform [4]–[6]. Hence, we compare our proposed algorithms against *CF*, the algorithm in [8] (called *LC*), and the two better measurement-based algorithms *Lo-U* and *No-U* in [9]. First, we consider a saturated network where all APs are transmitting downlink traffic. We set the number of orthogonal channels (K) to 3 to represent 802.11b/g; other larger values of K produce very similar trends as to those shown in Figs. 2-4, making our approach applicable to cellular networks and 802.11a. We consider three network sizes, three levels of rogue interference, and two network topologies, and thus have eighteen combinations ($3 \times 3 \times 2$), as shown in the x-axis of Fig. 3. The three network sizes include a 4-by-4 AP layout with 64 users, a 7-by-7 layout with 196 users, and a 10-by-10 layout with 400 users. Each AP may be associated with a different number of users; the average number of users for each AP is four. We consider low, medium, and high interference

from rogue interferers, where the ratio of the number of rogue interferers to the number of APs is 10%, 40%, and 70%, respectively. We consider a uniform topology where APs are regularly located as illustrated in Fig. 1(a), and a nonuniform topology, where APs are perturbed from the uniform layout with a small random distance (up to 25% of separation), as shown in Fig. 1(b). The separation between adjacent APs is 240 meters. We set q as 2. The noise floor is 10 dB above the thermal noise to properly represent the real world [12]. We assume each AP can source a max of 54 Mbps per the 802.11g standard.

For example, Fig. 2 considers 100 controlled APs with 10 rogues. *LC* in [8] is known to be the best prior to our previous work in [9]. Fig. 2 shows that *SS-R* and *SS-S* outperform *LC* by 13.1% and 16.8% in terms of mean user throughput, 13.6% and 18.5% in terms of median, 87.1% and 97.6% in terms of 25-percentile, and 1110% and 1180% in terms of 15-percentile user throughputs. *Lo-U* is known to be the best algorithm among all prior work, especially at uplifting throughputs for users with lowest throughputs. Nevertheless, Fig. 2 shows that *SS-S* outperforms *Lo-U* by 3.68%, 8.95%, 13.6%, 15.1%,

25.8%, and 72.6%, and *SS-R* outperforms *Lo-U* by -0.619%, 3.13%, 7.68%, 8.77%, 19.5%, and 84.9%, in terms of 50, 25, 20, 15, 10, and 5 percentiles of user throughputs, respectively. *SS-R* yields the highest 3-percentile throughput. Generally, *SS-R* sacrifices the users with higher throughput to improve the users with lower throughput. Although *SS-R* is worse than *SS-S* for users with high throughput, *SS-R* is still better than *Lo-U*, the best algorithm in the literature. Our algorithms yield enormous throughput gains especially for users with low throughputs. Fig. 3 shows that our algorithms enable more users to operate above 512 kbps irrespective of the number of APs and rogues; this trend is true for other throughput thresholds, as well. Fig. 3 shows that *SS-R* accommodates up to 18% and 7% more users than *LC* and *Lo-U*, respectively.

Above we assumed all traffic was *downlink* and optimized the frequency allocation for the *most active case* where *all APs are transmitting downlink traffic*. It is reasonable to optimize frequency allocation for this case, since in this case, frequency allocation is most crucial for interference mitigation. Now we examine the performance of the optimized frequency allocations in the presence of both downlink and uplink traffic. Work in [15] shows that uplink and downlink capacities in multiple cells are mutually coupled due to inter-cell interference, and no system-level analytic model has been found to model activities of multiple APs. In our next simulation, we consider that time is slotted, and propose an approximate probabilistic model where APs independently choose one of the three possible activity states at each time slot. An AP can be transmitting downlink traffic, receiving uplink traffic, or idle, with probabilities p_d, p_u, and $p_i = 1 - p_d - p_u$, respectively. For any AP that is transmitting downlink traffic or receiving uplink traffic at a certain time slot, a user is randomly chosen (with uniform probability distribution) out of all the users associated with this AP to be the recipient or the sender of the traffic. We fix the ratio of p_d to p_u as 5:1 [3], and and simulate 5 cases where $p_d + p_u$ (the probability that an AP is active) is 0.2, 0.4, ..., 1.0, respectively. We intend to see the effect of $p_d + p_u$ on the performance of the proposed algorithms. The assumption that the activity of each AP is independent from the other APs simplifies the simulations and provides a rule-of-thumb for the comparison. Fig. 4 shows that our algorithms consistently yield throughput gains (including both downlink and uplink) irrespective of the probability of AP activity; particularly the gains are high (up to 71% for 25-percentile throughput and 19% for median throughput) when APs are highly active (i.e., when the network traffic load is heavy). In Fig. 4 we see the same trend as in Fig. 2 that *SS-S* and *SS-R* are the best at providing high throughputs for users who suffer the lowest throughputs.

V. CONCLUSIONS

A central network controller with site specific knowledge can predict the path loss between any AP and client, and therefore predict the impact of SINR and throughput on every AP and user when the channel of any AP is changed. This site specific knowledge leads to vast network improvements which we have demonstrated by using two site-specific algorithms that can incorporate the importance of fairness across users. These algorithms substantially outperform all other published ones in various scenarios. Our proposed algorithms are particularly useful when the traffic load of the network is high and APs are highly active. The two algorithms, *SS-S* and *SS-R*, are better at uplifting the throughputs of users that suffer lowest throughputs when particular utility functions are chosen. Judicious selection of utility function is a topic of future research. We believe that site specific knowledge is also useful for other wireless communication problems in both cellular networks and WLANs, which will be validated by ongoing and future work; for example, work in [16] use site specific knowledge for load balancing in wireless networks.

REFERENCES

[1] S. Y. Seidel and T. S. Rappaport, "914 MHz path loss prediction models for indoor wireless communications in multifloored buildings," *IEEE Trans. Antennas Propag.*, vol. 40, no. 2, pp. 207 – 217, Feb. 1992.

[2] R. R. Skidmore, T. S. Rappaport, and A. L. Abbott, "Interactive coverage region and system design simulation for wireless communication systems in multifloored indoor environments: SMT Plus," in *Proc. IEEE International Conference on Universal Personal Communications*, vol. 2, Oct. 1996, pp. 646 – 650.

[3] C. Na, J. K. Chen, and T. S. Rappaport, "Measured traffic statistics and throughput of IEEE 802.11b public WLAN hotspots with three different applications," *IEEE Trans. Wireless Commun.*, vol. 5, no. 11, pp. 3296 – 3305, Nov. 2006.

[4] K. K. Leung and B.-J. Kim, "Frequency assignment for IEEE 802.11 wireless networks," in *Proceedings of IEEE Vehicular Technology Conference*, vol. 3, Oct. 2003, pp. 1422 – 1426.

[5] Y. Lee, K. Kim, and Y. Choi, "Optimization of AP placement and channel assignment in wireless LANs," in *Proc. IEEE Conf. on Local Computer Networks (LCN)*, Nov. 2002, pp. 831 – 836.

[6] A. Mishra, S. Banerjee, and W. Arbaugh, "Weighted coloring based channel assignment for WLANs," *ACM Mobile Computing and Communications Review*, vol. 9, no. 3, pp. 19 – 31, 2005.

[7] A. Mishra, V. Brik, S. Banerjee, A. Srinivasan, and W. Arbaugh, "A client-driven approach for channel management in wireless LANs," in *Proceedings of IEEE INFOCOM*, Apr. 2006, pp. 1 – 12.

[8] D. J. Leith and P. Clifford, "A self-managed distributed channel selection algorithm for WLANs," in *Proceedings of International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks*, Apr. 2006, pp. 1 – 9.

[9] J. K. Chen, G. de Veciana, and T. S. Rappaport, "Improved measurement-based frequency allocation algorithms for wireless networks," submitted to *IEEE Globecom*, Washington DC, USA, Nov. 2007.

[10] R. R. Skidmore and T. S. Rappaport, "System and method for measuring and monitoring wireless network performance in campus and indoor environments," U.S. Patent no. 7,096,160, August 22, 2006.

[11] P. Bahl and V. N. Padmanabhan, "A software system for locating mobile users: Design, evaluation, and lessons," Microsoft, Tech. Rep., Apr. 2000.

[12] J. K. Chen, T. S. Rappaport, and G. de Veciana, "Frequency allocation, transmit power control, and load balancing with site specific knowledge for optimizing wireless network performance," PhD Dissertation, University of Texas at Austin, May 2007.

[13] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556 – 567, Oct. 2000.

[14] S. Shakkottai, T. S. Rappaport, and P. C. Karlsson, "Cross-layer design for wireless networks," *IEEE Commun. Mag.*, vol. 41, no. 10, pp. 74 – 80, Oct. 2003.

[15] T. Bonald, S. Borst, N. Hegde, and A. Proutière, "Wireless data performance in multi-cell scenarios," in *Proceedings of ACM SIGMETRICS*, Jun. 2004, pp. 378 – 387.

[16] J. K. Chen, T. S. Rappaport, and G. de Veciana, "Iterative water-filling for load-balancing in wireless LAN or microcellular networks," in *IEEE Vehicular Technology Conference*, vol. 1, May 2006, pp. 117 – 121.