A DISTRIBUTED DYNAMIC CHANNEL ALLOCATION TECHNIQUE FOR THROUGHPUT IMPROVEMENT IN A DENSE WLAN ENVIRONMENT

Hui Luo, N. K. Shankaranarayanan

AT&T Labs – Research
200 S. Laurel Ave, Middletown, NJ 07748, USA

ABSTRACT

In a dense WLAN (WLAN) environment, the signal coverage area of each Access Point (AP) typically has significant overlap with that of the neighboring APs. This is a problem if there are limited frequency channels. This paper presents a distributed dynamic channel allocation algorithm that can improve per-user throughput significantly in dense WLANs, particularly for non-uniform traffic conditions. It is based on a cellular neural network model. Like a cellular neuron changing its state based on the information of its neighboring neurons, every AP determines the best channel it should use in the next time slot based solely on the traffic load of its neighboring APs and the channels used by them in the current time slot, but it actually switches to that channel with some fixed probability less than one. All APs in the network perform the above operation simultaneously. Computer simulations show that (1) given any traffic load distribution and any initial channel allocation, the algorithm converges to an equilibrium state in a short time, in which the overall throughput of the network is significantly improved; and (2) there exists an optimal switching probability that can minimize the time for the algorithm to reach the equilibrium state. The proposed technique has significant practical value due to its simplicity and effectiveness.

1. INTRODUCTION

IEEE 802.11 [1] is a hugely popular international standard for wireless LAN (WLAN) technologies that enable computers to communicate with each other or with a wired LAN over a distance of up to about 300 feet using a wireless channel on the 2.4 GHz or 5 GHz unlicensed band. The standard consists of MAC (medium access control) and PHY (physical layer) specifications. At present, IEEE 802.11 has three high-speed PHY specifications: 802.11b, 802.11a, and 802.11g. 802.11b supports a data rate of up to 11Mbps. Although 802.11b has 11 channels in the 2.4 GHz unlicensed band, there are only 3 non-overlapping channels. 802.11a supports a data rate of up to 54 Mbps; it specifies 12 non-overlapping channel on the 5 GHz unlicensed band. 802.11g supports the data rate as high as 802.11a, but remains compatible with 802.11b and it also has only 3 non-overlapping channels. Among these options, 802.11b has the largest installed base. It is predicted that 802.11g will prevail because of its compatibility with the popular 802.11b. The limited number of usable channels may have a significant impact on throughput in 802.11b/g dense WLANs. This is the motivation for our work.

Dense WLAN deployments are inevitable; (1) in order to eliminate coverage holes for a large-scale WLAN that covers an entire building or campus, and in order to maintain a high SNR to assure high data rates everywhere, the coverage areas of APs have to be overlapped; (2) in a crowded place, many APs with different owners are deployed without coordination. In these scenarios, if channels are inadequately allocated to neighboring APs, the throughput performance of WLAN stations associated with these APs will suffer, because all of them may have to compete for the same channel in order to exchange data with their APs, and thus making the channel overloaded. In the first scenario, the problem could be solved in the initial installation by a careful frequency planning and power control, but either the problem or coverage holes may appear in the future as local environment changes. In the second scenario, which will be very popular as more and more APs are installed by individuals, currently there are some simple solutions that address the problem only partially; some AP vendors build in a feature where an AP scans for radio energy and picks a channel with the least level, thus attempting to separate APs with the same channel. This is usually done when the AP is initialized; it cannot adapt to traffic dynamics afterwards.

This paper proposes a distributed dynamic channel allocation technique that optimizes the throughput of dense WLANs with strong cell overlap. It is based on a cellular neural network model. Like a cellular neuron changing its state based on the information of neighboring neurons, every AP determines the best channel it should use in the next time slot based solely on the traffic load of its neighboring APs and the channels used by them in the current time slot, but it actually switches to that channel at a fixed probability. All APs in the network perform the above operation simultaneously and repeatedly until they converge, i.e. they find their best channels. At this point, the per-user throughput is significantly improved throughout the WLAN. This technique is adaptive and distributed. It can be built into APs and will solve the problem for both the problem scenarios mentioned above.

The rest of paper is organized as follows. In Section II, the IEEE 802.11 MAC and its capacity analysis in the literature are briefly described as background. In Section III, the distributed dynamic channel allocation algorithm is described in detail. In Section IV, computer simulations are presented to demonstrate the effectiveness of the proposed method.

2. BACKGROUND

2.1. IEEE 802.11 MAC Specification

The IEEE 802.11 MAC employs CSMA/CA (carrier sense multiple access with collision avoidance) as the basic medium access mechanism to assure fair sharing of a WLAN channel by stations. When a station has a frame to transmit, its MAC layer waits for a required short time, then checks whether the channel is used or reserved by another station. If the channel is free, it transmits the frame immediately; otherwise, the MAC layer enters a deferred period that is randomly chosen according to its current contention window size. The deferred time elapses whenever the channel is free. At the end of the deferred period, the MAC layer transmits the frame. If the expected acknowledgement frame is not received on time, which means the receiving station does not receive the frame correctly, the MAC layer of the sending station will retry the transmission by following the above procedure with a doubled or maximum contention window size. If the transmission succeeds, the contention window size is reset to the minimum value.

Since an 802.11 WLAN operates in unlicensed spectrum, the wireless medium could be very noisy due to the existence of other types of radio devices operating on the same band and not conforming CSMA/CA, and thus data frames can be frequently corrupted. In order to assure reliable delivery of data frames over the noisy wireless medium, the IEEE 802.11 MAC employs a data frame exchange protocol for transmission of every data frame. It has a two-message form and a four-message form. The two-message form is suitable for short data frame transmissions. It works as follows: (1) Station A sends a data frame to Station B; (2) if it is correctly received, Station B sends a short ACK (acknowledge) frame to Station A; and (3) if Station A does not receive the ACK frame on time, it tries to re-send the data frame unless its MAC layer decides to drop the data frame after a sufficient number of retries. The four-message form is designed to overcome the "hidden node" problem. A "hidden node" is a station that is outside the sending station's radio range but inside the receiving station's radio range. Even if both the sending station and the hidden node run CSMA/CA, their data frames could collide at the receiving station, because neither of them can detect the other's signal or time window reservation announcement. The four-message data frame exchange protocol works as follows: (1) Station A sends a short RTS (request to send) frame to Station B, which indicates the time window that must be reserved for the forthcoming data frame transmission; (2) Station B sends a short CTS (clear to send) frame to Station A, which announces the reserved time window to all stations in Station B's radio range, including nodes that are hidden to Station A, and prevent them from sending during that period; (3) Station A sends the data frame to Station B in the reserved time window; and (4) if it is correctly received, Station B sends a short ACK frame to Station A.

2.2. IEEE 802.11 MAC Capacity

Many analytical and simulation results on IEEE 802.11 MAC capacity exist in the literature [2, 3]. In our work, we adopted the analytical model in [3] to quantify the relation of a single channel throughput vs. the number of WLAN stations (or load) using the channel. The model has been validated by simulations in [3]. Its assumptions closely match the real WLAN environment. That is, short frames are transmitted using the basic access method (CSMA/CA with Acknowledgement); long frames are transmitted using the RTS/CTS protocol; and both transmission methods may exist in the same WLAN.

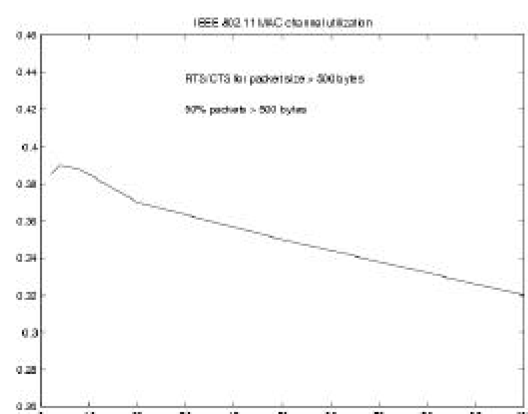

Fig. 1 IEEE 802.11 MAC throughput vs. traffic load.

The formula of the analytical model given in [3] is very complicated and thus omitted here. The result in [3] that will be used in this paper is plotted in Fig. 1. It shows the relation between the normalized throughput of a channel and the number of WLAN stations using the channel with the following assumptions: (1) any frame with more than 500 bytes is transmitted using the RTS/CTS protocol; and (2) the percentage of frames with less than 500 bytes is 50%. We use a function $P=f(M)$ to denote this relation, where M stands for the number of WLAN stations and P for the normalized channel throughput. In general, M denotes the load.

3. ALGORITHM

The distributed dynamic channel allocation algorithm is based on the following assumptions.

1. There is a dense wireless WLAN with multiple APs and stations. The signal coverage areas of neighboring APs are overlapped. Every AP has a neighborhood consisting of the coverage areas of neighborin APs.
2. Every AP periodically broadcasts how many WLAN stations (or the load) are associating with it. Every AP also detects the broadcast information from all of the APs in its neighborhood.
3. There are a set of non-overlapping channels in use. The maximum throughput of a channel in an AP's coverage area is determined by the total number of WLAN stations in the neighborhood area that are using that particular channel. The throughput of the channel is governed by the above $P=f(M)$ relation.

The objective hereby is to improve the overall throughput in the entire wireless network by improving the per-user throughput with every AP. This is done by optimal assignment of channels to APs. The algorithm is carried out by all AP simultaneously. We assume a rectangular array of APs where each AP is indexed by (i, j). Every AP performs the following operations periodically.

1. Compute its maximum throughput $P_k(i, j)$ if it uses channel k, for k = 1, 2, and 3, respectively (for 802.11b/g).

$$P_k(i,j) = \frac{M(i,j)}{\sum_{(m,n)\in S_k} M(m,n)} f\Big(\sum_{(m,n)\in S_k} M(m,n)\Big)$$

where, (i, j) denotes the position of an AP; $M(i,j)$ is the number of WLAN stations associating with the AP (i,j). S_k stands for the set of APs that are in the neighborhood of (i,j) and use the same channel k.

2. Find the best channel k_m such that the local throughput for the AP (i,j) is maximized.

$$P_{k_m}(i,j) = Max\{P_k(i,j) \mid k = 1,2,3\}$$

3. Switch to the channel k_m with a probability p, called the switching probability.

The APs do not need to stop the algorithm so that they can adapt changes in the wireless network. If no change happens, the algorithm will reach an equilibrium state quickly, in which every AP uses its best channel. We believe that the algorithm is very self-adaptive and should be insensitive to inaccuracies in the $P=f(M)$ equation. The analytical part of this algorithm is omitted due to paper length limitation. The effectiveness and convergence performance of the algorithm are demonstrated using computer simulations.

4. COMPUTER SIMULATIONS

We assumed a 10x10 square grid of 100 APs for our computer simulations. Every 3x3 grid of 9 APs formed a neighborhood for the AP in the center of that grid. We used 3 channels. Thus, some or all WLAN stations and APs in each neighborhood will have to share a channel. Note that if there are three APs X, Y, Z in a row that use the same channel, then the loads of X and Z will affect the throughput of Y, but the load of X will not affect the throughput of Z (and vice versa).

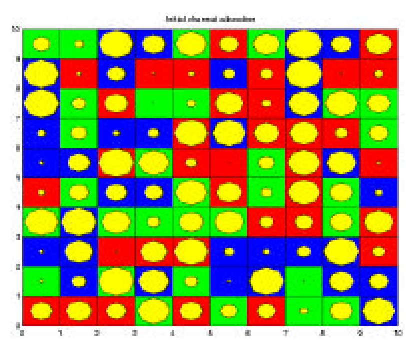

Fig. 2 Initial channel allocation of a 10x10 network

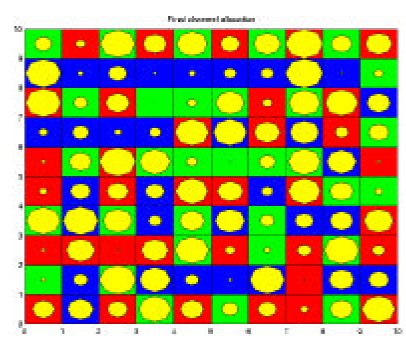

Fig. 3 Channel allocation at the equilibrium state

Fig. 2 shows an initial channel allocation and WLAN station distribution that are randomly generated. Every square stands for an AP and the WLAN stations associated with it. The number of WLAN stations (load) is represented by the size of the dot in the lattice. A larger dot means more WLAN stations. The color of the lattice (R/G/B) identifies the channel used by the AP and its stations. Fig. 3 shows the channel allocation after the algorithm converges to the equilibrium point. It can be seen that neighboring APs that have a large number of associated WLAN stations use different channels, and those having a few associating WLAN stations may share a channel. In other words, if an AP tends to be the largest dot in its grid, it will tend to get a channel (color) with minimal sharing, and adjacent APs with relatively light loads tend to be grouped together with the same channel (color).

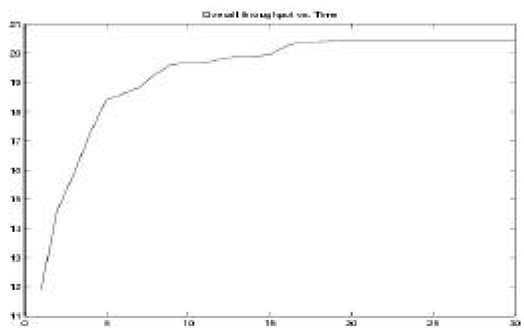

Fig. 4 Throughput improvement

Since in the equilibrium state every AP is using the best channel that maximizes its local throughput, the overall throughput of the dense WLAN should have been significantly improved. This is confirmed by the curve of overall throughput vs. running time plotted in Fig. 4, which shows a 70% improvement achieved by the algorithm in less than 20 iterations. This convergence speed is very fast for a 10 by 10 network, thanks to the distributed computing nature of the algorithm.

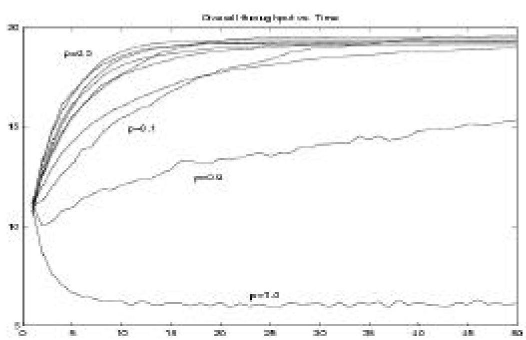

Fig. 5 Convergence speed vs. switching probability

Intensive computer simulations also reveal that, given a layout of a dense WLAN, there exists an optimal switching probability that can minimize the convergence time. Fig. 5 shows 10 curves, corresponding to a switching probability of 0.1, 0.2, …, 1.0, respectively. Each curve is an average of 1000 simulation runs with 1000 randomly generated initial channel allocation and load distribution for the 10 by 10 square layout. It can be seen that the switching probability 0.5 is the optimal one. Fig. 5 also shows that the algorithm won't work if there is no randomness in channel switch (i.e., the switching probability is 1.0).

5. CONCLUSIONS

As the IEEE 802.11 wireless is widely deployed, there are practical reasons that will lead to dense WLANs, which have significant overlap between the signal coverage areas of neighboring APs. Since the usable number of WLAN channels is very limited, WLAN stations may as a result suffer low throughput due to channel contention, especially when channels are inadequately allocated to neighboring APs. This paper presents a distributed dynamic channel allocation algorithm that can improve per-user throughput significantly in such scenarios. The algorithm is carried out by all APs in a parallel manner. Every AP determines the best channel it should use in the next time slot based solely on the number of WLAN stations associating with its neighboring APs and the channels used by them in the current time slot, and switches to that channel with a fixed probability. Computer simulations show that (1) given any traffic load distribution and any initial channel allocation, the algorithm converges to an equilibrium state in a short time, in which the overall throughput of the network is significantly improved; and (2) there exists an optimal switching probability that can minimize the time for the algorithm to reach equilibrium state.

REFERENCES

[1] ANSI/IEEE Std 802.11, Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 11: WLAN Medium Access Control (MAC) and Physical Layer (PHY) specifications, first edition, 1999.

[2] Giuseppe Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function", IEEE J-SAC, vol. 18, no. 3, March 2000, pp. 535-547.

[3] R. Bruno, M. Conti, and E. Gregori, "IEEE 802.11 Optimal Performances: RTS/CTS Mechanism vs. Basic Access", Proc. of PIMRC 2002.