2015 International Conference on Computational Science and Computational Intelligence

Real-Time Medical Data over WLANs

Constantine Coutras
Department of Computer Science
Montclair State University
Montclair, NJ 07043
Email: coutrasc@montclair.edu

***Abstract*—As medical records are moving on line, an increasing number of doctors are no longer keeping hand written notes. This finds them moving around in hospitals holding smart phones and tablets to enter and monitor medical data. In many cases it is very important to have any changes in medical data communicated immediately. While a variety of wireless systems are in place or have been proposed, they usually require doctors to carry around additional devices. In this paper the use the Point Coordination Function (PCF) of the IEEE 802.11 Wireless LAN protocol is proposed to transfer both typical and emergency medical data among doctors, databases and monitoring systems.**

***Index Terms*—Wireless Networks; IEEE 802.11; Medical Data.**

I. Introduction

Over the past few decades the increased level of public awareness concerning healthcare, physical activities, safety and environmental sensing has created an emerging need for smart technologies and monitoring devices able to sense, classify, and provide feedbacks to users' health status and physical activities, as well as to evaluate environmental and safety conditions in a pervasive, accurate and reliable fashion [1], [2], [3]. As medical records are moving on line, an increasing number of doctors are no longer keeping hand written notes. This finds them moving around in hospitals holding smart phones and tablets to enter and monitor medical data. While a variety of wireless systems are in place or have been proposed, they usually require doctors to carry around additional devices. In many cases it is very important to have any changes in medical data communicated immediately. The most prevalent wireless technology today is the IEEE 802.11 protocol for WLANs. Thus these WLANs can be used to transfer real-time medical data. The IEEE 802.11 MAC layer protocol provides *asynchronous, time-bounded, and contention free access control* on a variety of physical layers. The Point Coordination Function (PCF) provided by the IEEE 802.11 MAC protocol is designed to support time-bounded services, essential for transporting real-time services such as voice.

In this paper using the Point Coordination Function (PCF) of the IEEE 802.11 Wireless LAN protocol is proposed to communicate typical medical data, emergency medical data and even voice, to facilitate information exchange among doctors, databases and monitoring systems.

The paper is organized as follows: In Section 2 the IEEE 802.11 standard is presented in detail. Section 3 presents the proposal to support real-time medical services within the framework of the point coordination function. In Section 4 the specifics of the connection establishment procedure and numerical results are presented. Finally, in Section 5 areas of future research are presented.

II. Overview of the IEEE 802.11 Standard

An 802.11 network, in general, consists of Basic Service Sets (BSS) that are interconnected with a *Distribution System* (DS). Each BSS consists of mobile nodes, referred to as *stations*, that are controlled by a single *Coordination Function* — the logical function that determines when a station transmits and receives via the wireless medium. Stations in a BSS gain access to the DS and to stations in "remote" BSSs through an *Access Point* (AP). An AP is an entity that implements both the 802.11 and the DS MAC protocols.

The basic access method is the *Distributed Coordination Function* (DCF) which is known as *Carrier Sense Multiple Access with Collision Avoidance* (CSMA/CA). In addition to the DCF, the 802.11 also incorporates an alternative access method known as the *Point Coordination Function* (PCF) — an access method that is similar to "polling" and uses a point coordinator (usually the AP) to determine which station has the right to transmit. The PCF has been developed for providing real-time services. In this paper the focus is on the PCF.

978-1-4673-9795-7/15 $31.00 © 2015 IEEE
DOI 10.1109/CSCI.2015.145

A. Distributed Coordination Function

When using the DCF, a station, initially senses the channel to determine if another station is transmitting. The station proceeds with its transmission if the medium is determined to be idle for an interval that exceeds the *Distributed InterFrame Space* (DIFS). In case the medium is busy the transmission is deferred by the station until the end of the ongoing transmission. A random interval, referred to as the *backoff interval*, is then selected. The backoff timer is decremented only when the medium is idle; it is frozen when the medium is busy. Decrementing the backoff timer resumes only after the medium has been free longer than DIFS. A station can initiate transmission when the backoff timer reaches zero. To reduce the probability of collisions, after each unsuccessful transmission attempt, the backoff time is increased exponentially until a given maximum is reached.

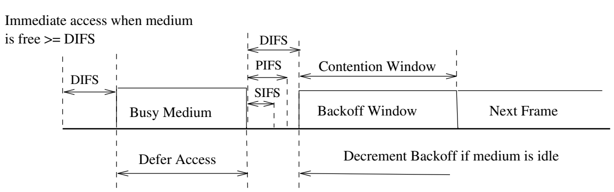

Fig. 1. Basic channel access method

Immediate positive acknowledgements are employed to determine the successive reception of each data frame. This is accomplished by allowing the receiver to transmit an acknowledgement after a time interval *Short InterFrame Space* (SIFS) (that is less than DIFS) immediately following the reception of the data frame. Acknowledgements are transmitted without the receiver sensing the state of the channel. In case an acknowledgement is not received the data frame is presumed lost and a retransmission of the data frame is scheduled. This access method, referred to as *Basic Access*, is summarized in Figure 1.

B. Point Coordination Function

The PCF is built using the DCF through the use of an access priority mechanism that provides synchronous or asynchronous data frames *contention free* access to the channel. In this case, contention and contention free periods alternate with each other as shown in Figure 2. A contention free period (during which PCF is active) and the following contention period (during which DCF is active) are together referred to as a *SuperFrame* (SF).

At the beginning of the nominal SF boundary the *Point Coordinator* (PC) senses the channel. If the channel is sensed to be idle, the PC seizes control of the channel by transmitting after it (the channel) has been idle for a time interval *Priority Interframe Space* (PIFS) that is chosen to be smaller than DIFS but larger than SIFS. However, if the medium is determined to be busy, the PC monitors the channel until it is idle, and then seizes its control by transmitting after the channel has been idle for PIFS. The PC maintains control of the channel throughout the contention free period by initiating transmissions after an idle period of PIFS. The transmission of a Contention Free Acknowledgement (CFACK) frame by the PC marks the end of the contention free period.

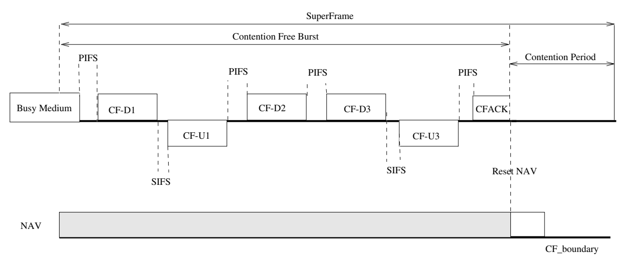

Fig. 2. Point Coordination Function

The PC sends data to stations in CF-Down frames which also achieve the polling function. A "poll bit," if enabled, in the CF-Down frame polls the destination of the CF-Down frame. A station transmits data in CF-Up frames, these CF-Up frames can be transmitted after the reception of a CF-Down frame with the poll bit enabled. The need for separate acknowledgements is avoided by "piggybacking" acknowledgements on subsequent frames, by the setting of an appropriate bit.

CF-Up frames are transmitted after the channel has been idle for a time interval SIFS as compared to CF-Down frames that are transmitted once the channel has been idle for PIFS. Thus the PC will transmit the next CF-Down frame in case there is no CF-Up frame in time interval PIFS after the transmission of the previous CF-Down frame. To minimize collisions during the contention free periods, each station sets its NAV equal to the maximum allowable length of the CF period. However, a station resets its NAV if a CF-ACK frame is seen by it before its NAV has expired. Refer to Figure 2 for more details. Station to station transfer of data frames is achieved by addressing the CF-Up frame to a destination station which then generates an acknowledgement following the rules of the basic access method; subsequently the PC seizes control of the channel.

The maximum length of the CF period is chosen such that at least one maximum sized 802.11 data frame, referred to as a MAC Protocol Data Unit (MPDU), can be transmitted in the contention period. Note that the length of a SF can differ from the nominal SF length due to a phenomenon called *superframe stretching*. Superframe stretching refers to the extension of a superframe beyond its nominal end time due to the ongoing transmission in the contention part of the SF (recall, the PC takes control of the channel only after the channel has been idle for a time interval of PIFS). The phenomenon of superframe stretching can lead to SF lengths being both smaller and larger than the nominal SF length. Note that it is the job of the PC to determine (i) appropriate "nominal" start times for the superframes, (ii) maintain polling lists of stations, and (iii) ensure that performance guarantees are met.

III. Support for Real-Time Services

Real-time medical data traffic demands strict performance guarantees from the network. Before a certain performance level can be guaranteed to a station, the characteristics of the traffic generated by it have to be known so that appropriate resources can be allocated to it. Therefore, stations that require performance guarantees are required to set up a *connection* with the PC. Thus, the term "connection" is used to refer to a contract between the PC and the station(s) where the PC guarantees the performance desired as long as the station does not violate the declared traffic characteristics.

The PC maintains a polling list that specifies the order in which stations are to be polled. The order in which the stations are polled is dynamic, i.e., it can be changed from one SF to the next. Typically, the polling list contains stations that have established connections with the PC. In addition, the PC *may* decide to poll stations that have not established connections. For simplicity the following assumptions are made:

- The nominal length and start times of each superframe (and therefore the nominal start times of the CF periods) is predetermined.
- The effect of hidden stations on the operation of the PCF is ignored.
- The contract associated with a connection cannot be renegotiated at any time.

In order to provide support for real-time services required by medical data it is essential to determine the time instants when any given station that has established a connection with the PC is to be polled.

Let C denote the raw transmission rate of the channel, l_{MPDU} the the length of an MPDU in bits and t_{MPDU} the *time* it takes to transmit an MPDU of length l_{MPDU} bits. Let $t(n), n = 1, 2, \ldots$, denote the nominal start time of the nth superframe and let $t'(n)$ denote the actual start time of the nth superframe. Further, t_{SF} is defined to be the nominal length of a superframe ($t_{\mathrm{SF}} = t(n+1) - t(n)$) and $t_{\mathrm{SF}}(n)$ and $t_{\mathrm{CF}}(n)$ denote the length of the nth superframe and contention free interval, respectively.

A polling list consists of stations transmitting synchronous and asynchronous data frames. Stations transporting synchronous frames should be given priority over stations transferring asynchronous data frames since no guarantees are provided to them. Within the set of synchronous stations the order in which the stations are polled should be such that the performance requirements of each are met even in the presence of both arriving and departing connection requests, and statistical variations in the amount of data transferred by various stations.

Throughout this paper the performance requirements are expressed in the form of a 2-tuple (D, ϵ). A performance requirement of (D, ϵ) indicates that no more than ϵ fraction of the traffic should exceed a delay of D time units.

To achieve the above performance objectives it is proposed that the PC be provided with the ability to guarantee each connection request it accepts "time window(s)" in each superframe within which it will be polled, and thus allowed to *initiate* the transfer of a data frame. In order to provide strict delay and loss guarantees it is important that (i) the position of this time window (relative to the nominal starting time of the SF) not change from one superframe to the next and (ii) the length of the time window be as small as possible. Intuitively, minimizing the tail distribution of the difference between nominal and actual poll times of any given station will lead to smaller delays (or smaller tail delay distributions). This is important since smaller tail delay distributions leads to more efficiency, especially for real-time traffic. To meet the requirements of real-time traffic it is *proposed* that

(P1) Each station be polled only once during each superframe.

(P2) Following a CF-Down frame each synchronous station (having established a connection) will send a single frame. The maximum frame size that can be used by a station should be negotiated at the time of connection establishment.

(P3) An arriving connection request, if accepted, is placed at the end of the polling list.

(P4) Following the departure of a connection request, the time window for each station in the polling list that is polled after the departing connection request, is advanced by the time allocated to the departing connection request in each superframe.

The following observations are important.

1) A window size of t_{MPDU} is the smallest that can be achieved and guaranteed.
2) The PC can advance the position of the time window for any given station in a SF if it promises to maintain the new position (with an acceptable variance) in all of the successive superframes. Therefore, the cumulative transmission time available to any station, after the service time window has been advanced, is at least as large as would have been made available to it had the time window not been advanced.
3) For any given station A, the set of stations that are polled after it, in any given SF, have no effect on the instant at which A initiates transmission to outside the service time window, in the SF under consideration.

The time at which a connection is polled in any given SF depends on (i) the start time of the SF, (ii) the size of the data frames transferred by stations polled before the connection of interest, and (iii) the arrival and departure of new connections to the polling list. Each of the above are addressed individually:

1) *Statistical Variations*: Consider any station i and any given SF; if all stations polled before station i in the SF fully utilize their allocated transmission times, the starting service time for station i relative to the starting time of the superframe does not change and therefore the polling time for station i lies within its preassigned window. However, if some of the stations that were polled before station i do not use the entire service time allocated to them, station i could be polled before its preassigned time window. In this case it is proposed that the PC poll an asynchronous station until the next polling instant lies in the service time window of station i. Since an asynchronous station is allowed to transmit for no more than t_{MPDU} each time it is polled, a service time window of length t_{MPDU} suffices. If no asynchronous station is ready to transmit then station i can be polled twice.
2) *Arriving Connection*: An arriving station, if it can be accepted, is put at the end of the polling list. Thus its addition does not impact the position of the service time window for any existing connection.
3) *Departing Connection*: Consider station i departing from the system. In this case the relative position of the service time window for all stations that follow station i in the polling list is *decremented* by the service time for station i in each SF. This has little impact on the service received by all existing connections in the polling list.

IV. Connection Establishment

Before a connection can be established for a given station it has to be ensured that the required transmission time can indeed be made available to the requesting station in the CF period of the superframe. Recall, that the maximum allowable length of the CF period is such that at least one maximum sized data frame can be transmitted during the following contention period. However $\mathrm{CF}^* := \mathrm{SF} - 2\mathrm{MPDU}$ is the maximum CF length that can be used to guarantee transmission time to stations transmitting synchronous data frames. Therefore, the sum of the allowable transmission times of all (synchronous) stations in the polling list should not exceed CF^*. Observe that the PC *may* choose a maximum contention free interval length that is smaller than CF^* in order to allow more time for the transfer of asynchronous data frames during the contention period.

The procedures to determine the service time required in each SF to meet the demands of a connection request are now presented . It is instructive to look at the availability of the channel from the perspective of a station that is allowed to transmit l^* bits each time it is polled. Only variable bit rate medical data are considered, as continuous bit rate data are not expected.

A. Variable Bit Rate Services

From the point of view of a single station the channel can be modeled as one that alternates between two states – "Off" and "On" in which it has a capacity of 0 and C, respectively. Under the proposals put forth in this paper the maximum time spent in the On and Off states is fixed (note that the exact time spent in each state is a random variable). For analytical tractability, the channel is characterized by an N-state Markov Modulated Fluid (MMF) source. Let C_n denote the channel capacity when the channel is in state $n, n = 1, \ldots, N$. Let $\beta_{ij}, i \neq j, i, j = 1, \ldots, N$, be the rate at which the channel moves from state i to state j; further define $\beta_{ii} := -\sum_{j=1, j\neq i}^{N} \beta_{ij}$.

Before the specific choice of parameters (the number of states and the transition rates between states) that enable us to model the channel with sufficient accuracy, is discussed, it is necessary to consider how to determine the buffer occupancy and delay distributions when both the arrival and service processes are modeled by general Markov modulated fluid sources.

Consider a station that generates traffic according to an M state MMF source. The arrival process is modeled by an MMF source since many real-time traffic sources, such as voice and video, have been characterized by Markov modulated fluid processes in the literature [4], [5], [6], [7].

Let λ_m be the rate at which traffic is generated when the station is in state $m, m = 1, \ldots, M$, and let $\alpha_{ij}, i \neq j, i, j = 1, \ldots, M$, be the rate at which the station moves from state i to state j; further define $\alpha_{ii} = -\sum_{j=1, j\neq i}^{M} \alpha_{ij}$.

Let $P_{mn}(t, x)$, denote the probability that at time t: (i) the station is in state m, (ii) the channel is in state n, and (iii) the station buffer contents are $\leq x$.

Let $F_{mn}(x)$ denote the buffer occupancy distribution in steady state, i.e., $F_{mn}(x) = \lim_{t\to\infty} P_{mn}(t, x)$.

And let $D(x)$ be a random variable that denotes the time spent in the system by a fluid particle that arrives when the amount of fluid in the buffer is x.

From [8] the probability that the time spent (a random variable) in the system by a fluid particle exceeds t is given to be:

$$P\{D \geq t\} = 1 - \sum_{n=1}^{N} \sum_{m=1}^{M} (F_{mn}(T) - \int_0^T P\{D_n(x) \geq t\} dF_{mn}(x)). \qquad (1)$$

B. Numerical Results

Numerical results follow. The interest is in the probability of overflow for various maximum delay times. This will determine the system's performance. Simulation results and theoretical calculations were so close, that the simulation results are not explicitly shown.

To evaluate the proposed operation of the IEEE 802.11 for real time medical data, it is evaluated under the most demanding conditions, which correspond to a real-time voice. Consider a 32 Kbps voice source that has access to a 1Mbps channel, which polls the source according to the proposal. The source has an average On time of 40msec and an average Off time again of 40msec. The source is modeled by a 2 state (one On and one Off state) MMF source. The channel has a superframe (t_{SF}) of 40msec. It is modeled by a 16 state (8 On and 8 Off) MMF source. This is done because the channel is much more "deterministic" in its behavior (however, the existence of "superframe stretching" does not allow for the use a deterministic model). The On and Off times (T_{on} and T_{off}) for the channel are calculated so that the source gets the appropriate bandwidth. When the bandwidth allocation factor BA_{factor} is 1.0 the source is given 32kbits of bandwidth (this of course is twice its mean value, since the on and off periods are equal).

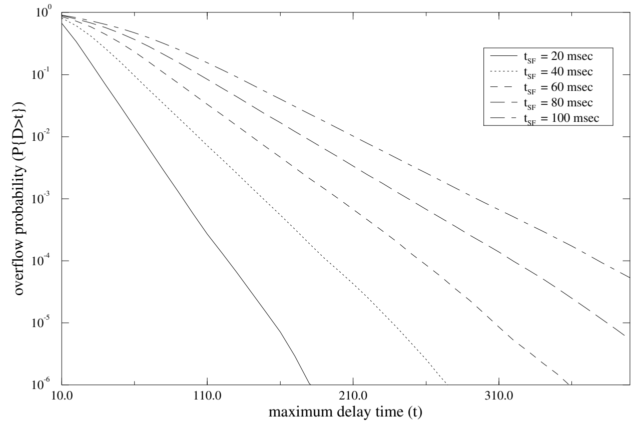

Fig. 3. Superframe variation

These are the default values. The graphs that vary some of these values state the values of that variation.

In Figure 3, the probability of overflow for various maximun delay time values is shown. Each curve represents a different channel superframe value. The smaller the superframe size (t_{SF}), the better the performance. A smaller superframe size actually means smaller intervals between polling times for the source. As observed, the superframe size (t_{SF}) is a very important parameter in determining the performance of the system.

In Figure 4, the probability of overflow is shown again, but now each curve is obtained by varying the BA_{factor} . Giving the source a BA_{factor} that is more than one is actually giving it more than 32kbps during each superframe. The performance improves as the BA_{factor} increases, although beyond a BA_{factor} of two, one observes diminishing returns for this case.

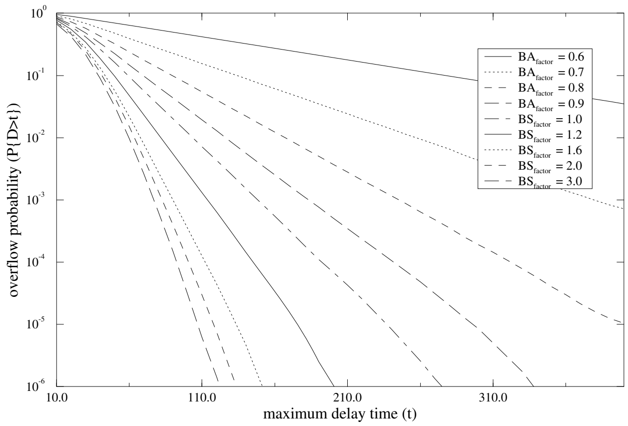

Fig. 4. Bandwidth variation

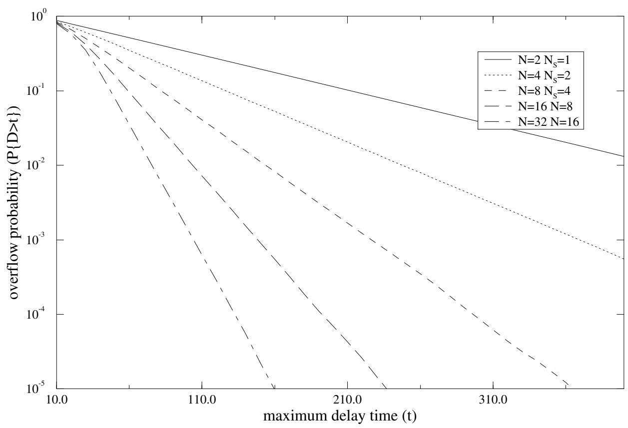

Fig. 5. Channel states variation

Giving the source a BA_{factor} that is less than one is actually giving it less than 32kbps during each superframe. As the BA_{factor} goes towards 0.5 (the bandwidth allocation equal to the mean value of sources rate) the performance drops quickly.

In Figure 5 the probability of overflow for various maximun delay time values is shown. Each curve represents a different number of states that model the channel. Using more states to represent the channel makes its behavior more "deterministic" and its superframe more "stable." So here the effect of the channel's superframe "stability" on overflow probability is seen. The more the channel states modeling the channel, the less the effect of the superframe stretching, and the better the performance.

Finally, in Figure 6 the probability of overflow for various maximun delay time values is shown again. Now each curve represents a source with a different pair of On and Off times. Observe that the smaller the On and Off periods the better the performance, as the source is less bursty.

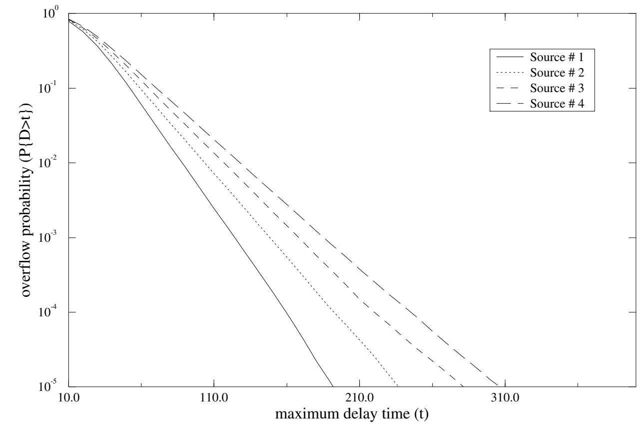

Fig. 6. Source variation

V. Conclusion

A brief outline of the IEEE 802.11 MAC protocol was presented and the procedures required for supporting real time medical data were discussed. The key engineering decisions that must be made in order to efficiently support these services were identified. Resource allocation and call admission/rejection in IEEE 802.11 WLANs was discussed. Future research will focus on software development to take advantage of the use of the IEEE 802.11 protocol as shown in this paper for real time medical data transfer.

References

[1] S. Oha, S. Ghoseb, K. C. andJ. Ryud, and J. Hane, "Recent trends in convergence-based smart healthcare service," *Technology and Health Care*, vol. 22, pp. 303–307, 2014.

[2] G. Ciuti, L. Ricotti, A. Menciassi, and P. Dario, "Omems sensor technologies for human centred applications in healthcare, physical activities, safety and environmental sensing: A review on research activities in italy," *Sensors - Open Access Journal*, vol. 15, pp. 6441–6468, 2015.

[3] S. K. Vashist, E. M. Schneider, and J. Luong, "Commercial smartphone-based devices and smart applications for personalized healthcare monitoring and management," *Diagnostics*, vol. 4, pp. 104–128, 2014.

[4] J. Daigle and J. Langford, "Models for analysis of packet voice communications systems," *IEEE Journal on Selected Areas in Communications*, vol. SAC-4, pp. 847–855, September 1986.

[5] A. Elwalid and D. Mitra, "Approximations and admission control of a multi-service multiplexing system with priorities," in *Proc. IEEE INFOCOM96*. Boston, Massachussetts, April 1995, pp. 463–472.

[6] A. I. Elwalid, "Markov modulated rate processes for modeling, analysis and control of communication networks," Ph.D. dissertation, Graduate School of Arts and Sciences, Columbia University, 1991.

[7] B. Maglaris, D. Anastassiou, P. Sen, G. Karlsson, and J. D. Robbins, "Performance models of statistical multiplexing in packet video communications," *IEEE Transactions on Communications*, vol. 36, no. 7, pp. 834–44, July 1988.

[8] C. Coutras, S. Gupta, and N. Shroff, "Scheduling of real time traffic over IEEE 802.11 WLANs," *Baltzer Journal of Wireless Networks*, vol. 6, pp. 457–466, 2000.