A Generic Framework for Modeling MAC Protocols in Wireless Broadband Access Networks

Xinhua Ling, Jon W. Mark, and Xuemin (Sherman) Shen, University of Waterloo
Yu Cheng, Illinois Institute of Technology

Abstract

In this article we present a simple yet accurate generic analytical model for a family of slotted CSMA/CA-based MAC protocols widely used in various wireless broadband access networks. The proposed model is based on a hierarchical three-level renewal process concept, which leads directly to some important MAC protocol performance metrics, such as throughput and average frame service time. The applicability of the model is demonstrated by extensive simulation results.

Wireless broadband access networks have become indispensable with the rapid growth of demand for wireless access to the Internet. The everlasting challenge of how to further improve resource utilization efficiency and provide better quality of service (QoS) attracts great efforts from both industry and academia. Medium access control (MAC) protocols play a critical role in determining the performance of wireless broadband access networks, which directly affects the perceived QoS of end users. The design and analysis of MAC protocols for such networks have gained much attention from researchers. In the former, a number of standards have been ratified for various types of broadband access networks such as wireless local area networks [1] and wireless metropolitan area networks [2]. In the latter, performance evaluation of these protocols has been carried out by both simulations and theoretical modeling approaches. While simulation studies, usually time consuming, may only address particular scenarios under specific conditions, analytical modeling enables one to gain deeper insight into the characteristics of the protocol.

The performance of MAC protocols is traditionally analyzed by developing stochastic models, often with various assumptions and approximations. In the literature, there are mainly three techniques commonly used in this area.

- *Traditional S-G analysis* [3] was widely used in the 1970s and '80s to analyze the throughput-delay performance of both slotted and non-slotted multiple access protocols such as ALOHA and carrier sense multiple access (CSMA). It assumes that an infinite number of nodes collectively generate traffic equivalent to an independent Poisson source with an aggregate mean packet generating rate of S packets/slot, and the aggregated new transmissions and retransmissions are approximated as a Poisson process with a rate of G packets/slot. The scenario considered is mainly of theoretical interest in the sense that a practical system has just a finite number of users, each of which usually has a buffer size larger than one, rather than that assumed in the S-G analysis.
- *Markov analysis* is another widely used technique. A Markovian model of the system is developed and its state transition probabilities need to be found. The state space of the Markovian model increases with both the complexity of the protocol studied and the number of users in the system, which hinders its usage in a system with a large user population.
- *Equilibrium point analysis* (EPA) is a fluid-type approximation analysis usually applied to systems in steady state [4]. It assumes that the system always works at its equilibrium point so that the number of users in any working mode is always fixed. To facilitate analysis, it requires a set of nonlinear equations, the number of which equals the number of working modes (e.g., different backoff stages) in the system.

A simple yet accurate analytical model is proposed in this article to analyze a family of CSMA with collision avoidance (CSMA/CA)-based MAC protocols commonly used in wireless broadband access networks. We model the behavior of an individual node instead of the channel. The proposed model is based on a novel concept of a *hierarchical three-level renewal process*, which can be solved by the fixed point technique. The new modeling approach significantly simplifies the mathematical analysis, where the important performance metrics of MAC throughput and average frame service time can be directly obtained. The proposed model is a general framework that is applicable to CSMA/CA-based MAC protocols with various backoff policies. A simple comparison among the proposed approach and the aforementioned three is given in Table 1.

In the remainder of this article we first describe the network considered, and then present the three-level renewal process and the MAC performance analysis based on it. Several popular MAC protocols are studied as applications of the proposed analytical model. Numerical results are given to verify the accuracy of this model. Concluding remarks are then given.

 0890-8044/07/$20.00

	S-G	Markov	EPA	Proposed
Station population	Infinite	Finite	Finite	Finite
Model complexity increases with backoff policy complexity	No	Yes	Yes	No
Handling unsaturated stations	Yes, limited to buffer size = 1	Yes, complexity increases with buffer size, or need help of queueing analysis	Yes, complexity increases with buffer size	Yes, no constraint to buffer size
Traffic model	Poisson	Poisson or Bernoulli	Bernoulli	General

■ Table 1. *Comparison of performance modeling approaches.*

Network Model

A single-hop wireless access network consisting of N functionally identical stations and a central receiver (e.g., an access point in a WLAN or a base station in a cellular system) is considered. Specifically, there are no hidden terminals in the network.[1] The time axis is slotted, and all the stations are synchronized so that all stations start their transmission only at the beginning of a slot. An ideal wireless channel without transmission error is assumed so that all transmitted frames may be lost only due to collisions caused by simultaneous transmissions from two or more nodes. However, imperfect channels that may cause unsuccessful reception of frames due to transmission error can also be embedded in our analysis, by taking into consideration the independent packet error probability[2] as in [6]. All MAC frames are assumed to have the same fixed length, which is a widely adopted assumption in MAC protocol analysis [3] and can easily be achieved in practice by commonly used link layer functions, such as fragmentation or concatenation of the upper layer packets. A short acknowledgment (ACK) frame is transmitted by the receiver immediately after every successful MAC frame transmissions, and a negative ACK (NACK) frame is transmitted in response to a collision. Alternatively, the sending stations will determine that there is a collision if the ACK frame is not received within a timeout period in which no station other than the receiver is allowed to transmit. In this case the timeout period can be treated as if it is occupied by a virtual NACK frame transmitted by the receiver. The aggregate transmission time of the MAC frame and the associated ACK or NACK is L slots, and no new transmission will start during this period.

In this article we focus on the uplink (from the stations to the central receiver), since this is where random access MAC protocols are usually used in current wireless broadband access networks. Also, we concentrate mainly on the *saturation case* in which every station always has at least one frame in the MAC buffer waiting for transmission. The throughput in this case, *saturation throughput*, is a fundamental performance figure defined as the limit reachable by the system throughput as the offered load increases, and represents the maximum load the system can carry in stable conditions [7]. How to develop the analytical model for networks with unsaturated stations is also discussed briefly.

[1] In a typical access network, all the stations can correctly sense the channel status, although they may not be able to correctly receive frames from all other stations [5].

[2] Since the channel fading process is independent of the packet transmission process, the events of packet errors and packet collisions are mutually independent.

Renewal Process Based Framework

In a steady state network with a CSMA/CA-based MAC protocol, a randomly tagged station contends for channel access following the same rule for each frame, i.e., it uses the same initial parameters pertaining to backoff, transmitting, and possible retransmitting processes for each frame. Therefore, the channel access process of an individual station is regenerative with respect to the time instants of the completion of each successful frame transmission. The period between two consecutive successful frame transmissions from the same station thus forms a renewal cycle in a renewal process [8], which directly relates to some important MAC performance metrics such as MAC throughput and average frame service time. This key observation inspires the basic analytical model presented below.

Three-Level Renewal Process

For many CSMA/CA-based MAC protocols, a hierarchical three-level renewal process, as illustrated in Fig. 1, can model precisely the behavior of a tagged station. A common characteristic of these protocols is that a certain form of backoff procedure is required before a station can access the channel, and the backoff procedure may be interrupted usually by transmissions from other stations and resumes after the channel becomes idle again. The time instants at which the backoff

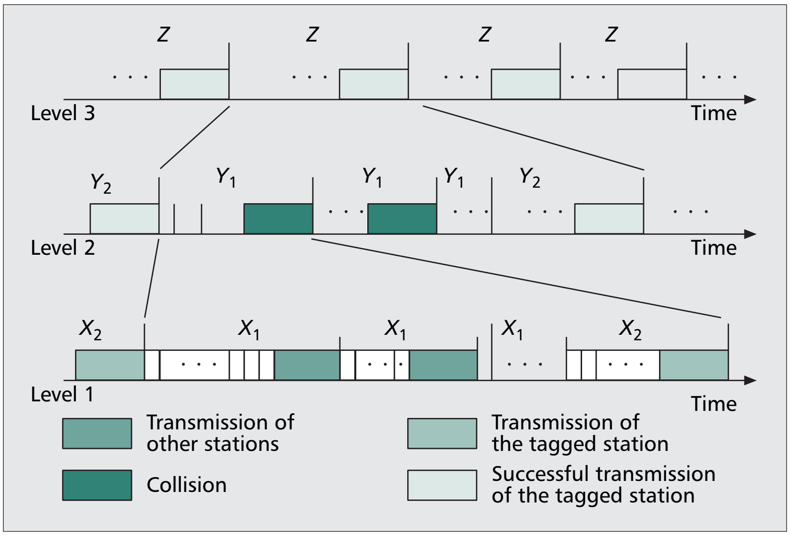

■ Figure 1. *Illustration of the concept of 3-level renewal process.*

procedure resumes (or restarts itself) can naturally be viewed as basic renewal points. Over a larger timescale, the end of each transmission from the tagged station is a higher level of renewal point of the frame service process. If the timescale is even larger, the renewal points can also be set at the end of each *successful* transmission from the tagged station, which delimits the important highest-level renewal cycle identified earlier.

In Fig. 1 a level 1 renewal cycle X is defined as the period between the end of a channel busy event and the end of the next one. There may be a number of idle slots in which no station transmits between two consecutive channel busy events. From the viewpoint of the tagged station, a level 1 cycle is of type X_1 if the busy channel is caused by transmission from other stations, or of type X_2 if its own transmission causes the channel busy. Notice that the transmission in an X_2 level 1 cycle may be a successful transmission or a collision.

A level 2 renewal cycle Y is from the end of an X_2 cycle to the end of the next X_2 cycle. As shown in Fig. 1, there can be a number of X_1 cycles before the X_2 cycle. Depending on the result of transmission in the X_2 cycle included, a level 2 cycle can be of type Y_1, in which the transmission results in a collision, or of type Y_2, in which the transmission succeeds.

Finally, a level 3 renewal cycle Z is from the end of a Y_2 level 2 cycle to the end of the next Y_2 cycle. Similarly, there can be a number of Y_1 cycles before the Y_2 cycle. Therefore, the successful transmission of a frame in the Z cycle can be viewed as the reward for the level 3 renewal cycle. The throughput of the tagged node can thus be obtained as the average reward in a Z cycle.

MAC Performance Analysis

As an important MAC performance metric, the average frame service time (or access delay) T_s is defined as the average duration from the instant that a frame becomes the head-of-line at the MAC buffer to the end of its successful transmission. As can be seen, T_s is exactly the average length of the level 3 renewal cycle in the proposed analytical model. The normalized throughput obtained by an individual station, per station MAC throughput, is the ratio of the transmission time of the MAC frame over T_s, and the saturation throughput of the network is thus simply the aggregation of the per station throughput.

To obtain the average length of the level 3 cycle, we need to find the characteristics of the random variables that determine the structure of the whole hierarchical renewal process. Depending on the specific channel access policy, the basic fixed point equations extracted from the analysis of level 1 cycles may be slightly different for each specific protocol. After the average length of a level-1 cycle $E[X]$ has been obtained, however, the analysis of the other two levels is shared by different protocols, which is one of the advantages of the proposed model.

As shown in Fig. 1, a level 2 cycle contains a number j of type X_1 cycles and an ending X_2 cycle, where j follows a geometric distribution with parameter P_{tx}, the probability that a level 1 cycle is of type X_2. Therefore, the average length of a level 2 cycle is $E[Y] = E[X]/P_{tx}$. Similarly, the number of level 2 cycles contained in a level 3 cycle also follows a geometric distribution, but with parameter P_{suc}, the probability that a transmission from the tagged station is successful. Thus, the average length of a level 3 cycle is

$$E[Z] = \frac{E[Y]}{P_{suc}}, \tag{1}$$

the normalized per station throughput is directly given by

$$\eta_s = \frac{L}{E[Y]}, \tag{2}$$

and the aggregated throughput is thus $N\eta_s$ for a homogeneous network.

Extension to Unsaturated Stations

The analysis so far has been for the saturated stations. In practice, this corresponds to the case when some delay-insensitive data applications (e.g., bulk file transfer through FTP) run on the stations. With multimedia applications that usually generate bursty traffic, a station may work in the unsaturated situation in which the MAC buffer may be empty from time to time. The proposed analytical model can also be extended to unsaturated stations. In this case, a station will contend to access the channel only when it has a frame in the buffer waiting for transmission, which occurs with the probability of having a nonempty buffer. This probability is equal to the ratio of the average frame service time over a given average frame inter-arrival time. The former can be obtained from the above model for saturated stations with slight changes, reflecting the effects of the aforementioned probability [9].

Application Instances

For CSMA/CA based protocols, the proposed hierarchical three-level renewal process based model makes the performance analysis easy to follow. The representative protocol is the legacy slotted p-persistent CSMA/CA [10]. Other good examples include the MAC protocol for the contention access periods of the IEEE 802.15.4 standard [11] and the IEEE 802.11 distributed coordination function (DCF), which is the de facto standard MAC protocol for WLANs. These three protocols are studied in the following to demonstrate the application of the proposed framework.

Legacy p-Persistent CSMA/CA

In the legacy slotted p-persistent CSMA/CA protocol, a station will sense the channel when it has a frame for transmission. If the channel is idle, the station transmits the frame with probability p. With probability $1 - p$, the station will defer its transmission decision by one slot. If the channel is still idle in the next slot, the station will repeat the above procedure. When the channel is sensed busy, the station waits until the channel becomes idle again and then operates as above. This probabilistic channel access rule can indicate that the station has adopted a geometric backoff policy, and the backoff procedure stops when there is a transmission from other stations and restarts itself after the transmissions ends. From this perspective, this protocol is an exact fit for the proposed three-level renewal process model.

The fixed transmitting probability p in an idle slot is the key to analyzing this protocol. From Fig. 1, a level 1 cycle contains a random number of idle slots, which follows a geometric distribution with parameter $(1 - p)^N$, and an ending block of channel busy slots with length L. The mean length of a level 1 cycle can thus be obtained straightforwardly, as can the average lengths of level 2 and level 3 cycles by following the approach outlined earlier.

IEEE 802.15.4 Contention Access

In the contention access period of the beacon-enabled mode specified in the IEEE 802.15.4 standard, each station adopts a slotted nonpersistent CSMA/CA mechanism with binary exponential backoff, termed *CAP-MAC*. According to this protocol, a station with a frame waiting for transmission in the MAC buffer is required to backoff a random number of slots first. The station does not sense the channel in the backoff slots, which is opposite to the *p*-persistent CSMA/CA protocol. At the end of this backoff stage, the station will perform the first channel clear assessment (CCA). If the channel is sensed idle, the station will conduct the second CCA in the next slot. Only when both CCAs indicate an idle channel will the station start the transmission in the next slot; otherwise, it will enter the next backoff stage. The number of backoff slots in stage m is uniformly distributed over $[0, 2^{BE_m})$, $0 \le m \le M$, where M is the maximum number of backoff stages allowed for a frame, BE_0 is the minimum and initial backoff exponent, and $BE_{m+1} = BE_m + 1$ and is upper-bounded. If all the M backoff stages end up with a busy channel indicated by the associated CCAs, a *channel access failure* will be reported. The node may then start the above procedure again for the next frame.

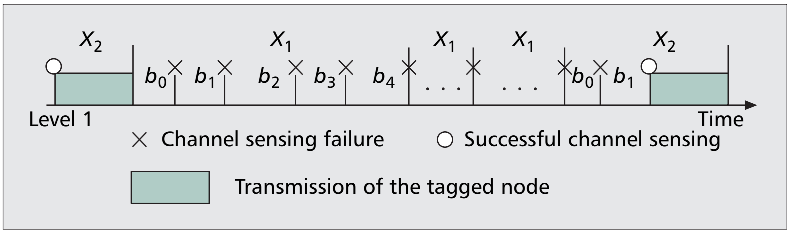

Figure 2. *The level 1 renewal cycle for IEEE 802.15.4 MAC.*

According to the above protocol, the renewal point of the level 1 renewal cycle can be set when a station resets its contention window to the minimum value after each transmission trial or when it senses a busy channel at the end of the Mth backoff stage. Correspondingly, in Fig. 2 a level 1 renewal cycle is defined as the period between the two adjacent time instants that the tagged node starts a stage 0 backoff. In this case, the X_1 type is a cycle that includes no transmission from the tagged node, resulting from M channel sensing failures; and the X_2 type is a cycle that contains a period L of transmission from the tagged node, after m backoff stages, $0 \le m \le M$. Again, the ending points of the X_2 cycles delimit the level-2 cycles as in the basic model. Thus, the level-2 and level-3 renewal cycles remain the same as in the basic model.

To analyze this variant of the model, the number of *channel sensing attempts V* conducted by the tagged node can be viewed as a reward associated to a level 1 renewal cycle. For a homogeneous network, the sensing failure probability α is the same for all the channel sensing activities from all the stations. With a given α, the number of sensing attempts for one station in a level 1 renewal cycle follows a truncated geometric distribution with parameters α and M; so does the number of backoff stages contained in a level 1 cycle. The average length of a level 1 cycle can thus be obtained easily [9]. According to the renewal reward theorem [8], the channel sensing probability β for the tagged station is given by the ratio of the average sensing attempts over the average length of the level 1 cycle, which is a function of α. On the other hand, the channel sensing failure probability α can be obtained as a function of β by carefully studying the transition probabilities among the channel states. The equations that express the relationship between α and β can be solved by the fixed point technique to obtain α and β, which can then be used to derive the average lengths of the three-level cycles and, finally, the desired MAC performance metrics [9].

Notice that if the same backoff policy is adopted, the legacy slotted *nonpersistent* CSMA/CA [10] differs from this protocol only in that the former requires just one successful channel sensing, while the latter needs two consecutive such successes to enable the frame transmission. Therefore, the above analysis can easily be adjusted for the legacy nonpersistent CSMA/CA protocol.

IEEE 802.11 Distributed Coordination Function

The IEEE 802.11 standard [1] distributed coordination function (DCF) specifies the channel access method as follows. When a station has a frame in the MAC sublayer buffer, it will first sense the channel. If the channel is busy, it will set its backoff counter (BC) to a random number uniformly distributed over $[0, CW_m]$, with CW_m the contention window (CW) for retransmission stage m, $0 \le m \le M$. After the channel becomes idle for a period of DCF-interframe-space (DIFS), the BC will count down for each time slot with idle channel, freeze when the channel is busy due to transmission(s) from other station(s), and resume after the channel turns idle for a period of DIFS again. When the BC reaches zero, the station will transmit immediately. A station will reset its CW to the minimum value CW_0 after each successful transmission. If two or more stations transmit at the same time slot, a collision occurs and all involved stations will double their *CWs* (upper-bounded by CW_{max}) and backoff again. The frame will be discarded if retransmission fails after a predefined retry limit, which depends on the frame length.

Applying the proposed framework to DCF, we again set a level 1 cycle as the period between two consecutive transmissions from the stations as in the basic model. The level 2 cycles are actually *semi-renewal* due to the changing backoff parameters. However, the renewal points for level 3 cycles remain to be the time instants when a new frame starts to be served. Notice that they are equivalent to the instants when the CW is reset to CW_0, which are the renewal points for level 1 cycles in the model for the CAP-MAC. This resemblance inspires us to analyze DCF directly at level 3, using an approach similar to the one for analyzing the level 1 cycles in CAP-MAC, outlined as follows.

The assumption that each transmission results in a collision with a constant probability γ regardless of the transmission history of the frame [7] is adopted in this analysis. With a given γ, the number of *transmission trials R* of a frame is thus a random variable with geometric distribution. For each new frame, R can also be deemed as a reward associated to the level 3 renewal cycle. Therefore, by the renewal reward theorem, the probability τ of a station to transmit (i.e., access the channel) at the beginning of a randomly chosen generic slot[3] is given by the ratio of the average number of transmission trials to the average total number of generic slots in the level 3 cycle. On the other hand, the collision probability can be obtained as a simple function of τ, considering that a collision occurs only when two or more stations transmit simultaneously. These resultant equations can be jointly solved to obtain the above two probabilities. The level 3 cycle can thus be deemed as a number of generic slots in which the tagged station either does not transmit or transmits but encounters a collision, followed by an ending generic slot containing the successful transmission from the tagged station. Consequently, the average length of the level 3 cycle in units of generic slots has a geometric distribution with parameter $\tau(1 - \gamma)$.

It is important to properly determine the *fixed point*, which is the key to the quantitative analysis in the proposed analytical framework. To model the MAC protocol, the criterion for selecting the fixed point is that the MAC behavior of each node can be independently modeled around the parameter

[3] *A generic slot may refer to an idle time slot, a successful transmission or a collision with respective probabilities, as defined in [7].*

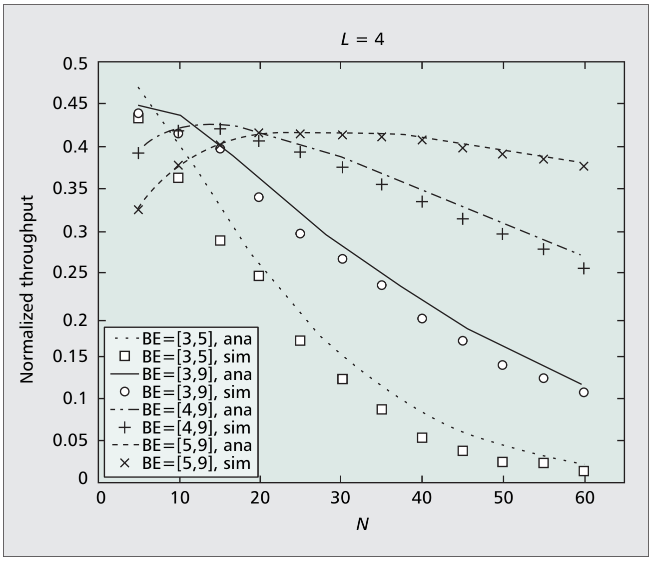

■ Figure 3. *Saturation throughput of the IEEE 802.15.4 CAP-MAC.*

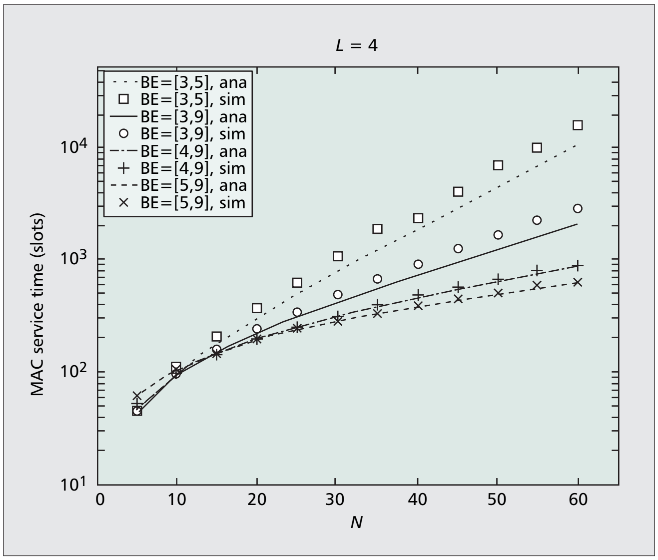

■ Figure 4. *Average frame service time in saturated case.*

associated with the fixed point; the equations describing different nodes are coupled by the fixed point. In analyzing the CAP-MAC, the probability to start sensing the channel is selected as the fixed point; in contrast, the channel access probability serves as the fixed point in the analysis of the *p*-persistent CSMA/CA and the IEEE 802.11 DCF. It is noteworthy that if an improper fixed point without the independency property is selected, the analytical model will lead to inaccurate performance results.

Numerical Results

In this section we use the IEEE 802.15.4 CAP-MAC as an example to assess the accuracy of the analytical model. Analytical results are compared with simulation ones for the MAC throughput in Fig. 3, and for average frame service time in Fig. 4, respectively. Results for several variants of the CAP-MAC with different parameter settings (e.g., different minimum and maximum backoff exponents) have been given to verify the versatility of the proposed model. Figures 3 and 4 demonstrate the high accuracy of the proposed model in all the scenarios.

To demonstrate the effectiveness of the model for the unsaturated case, Fig. 5 shows the average frame service time vs the frame arrival rate for a network with 20 stations. Poisson traffic is used in the simulation to compare with the analytical results. The average frame service time increases quite slowly with low to medium load, and it soars when the load becomes high until it reaches a saturation level which depends only on the station population of the network. It can be seen that the analytical results approximate the simulation ones very well.

Concluding Remarks

We have presented a simple yet accurate analytical model for a family of CSMA/CA based MAC protocols. By proposing a novel hierarchical three-level renewal process mechanism and applying the fixed point technique, we can describe and analyze the protocol in a straightforward manner. In each application instance of the model, a pair of fixed point equations is obtained by properly adjusting the parameters to capture the salient features of the backoff procedure and channel access policy in the protocol studied. After the fixed point equations are solved, the MAC throughput and the average frame service time for an individual station are obtained immediately,

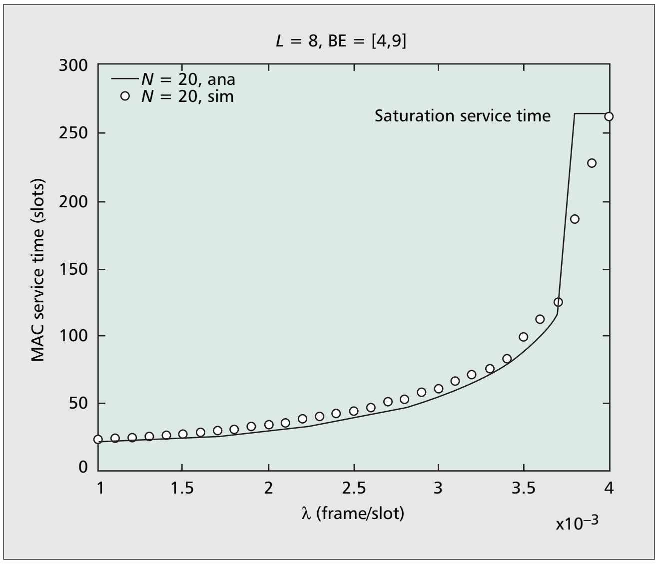

■ Figure 5. *Average frame service time in unsaturated case.*

due to the direct relationship between the level 3 renewal cycle and the frame service time. The accuracy of the analysis has been demonstrated by extensive simulation results.

The proposed model is in fact a very general analytical framework. It can be used for MAC protocols with various types of backoff policies (e.g., multiplicative increase instead of exponential increase, bounded or un-bounded), as long as the backoff procedure resets to its initial status for each new frame transmission. The modeling approach can also be extended to analyze networks with heterogeneous stations with different backoff parameters such as contention window and inter-frame spaces. One such application of the proposed model is a recent work of the performance analysis of the prioritized channel access MAC protocol in multiband orthogonal-frequency-division-multiplexing-based ultra wideband WPANs [12].

Acknowledgment

This work has been supported jointly by the Natural Sciences and Engineering Council (NSERC) of Canada under Strategic Grant # STPGP 257682 and Research In Motion (RIM).

References

[1] IEEE 802.11, "Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications," 1999.

[2] IEEE 802.16, "Local and Metropolitan Area Networks part 16: Air Interface for Fixed Broadband Wireless Access Systems, 2004.

[3] D. Bertsekas and R. Gallager, *Data Networks*, 2nd ed., Prentice Hall, 1992.

[4] S. B. Tasaka, *Performance Analysis of Multiple Access Protocols*, MIT Press, 1986.

[5] H. Zhai, X. Chen, and Y. Fang, "How Well Can the IEEE 802.11 Wireless LAN Support Quality of Service?," *IEEE Trans. Wireless Commun.*, vol. 4, no. 6, 2005, pp. 3084–94.

[6] X. J. Dong and P. Varaiya, "Saturation Throughput Analysis of IEEE 802.11 Wireless LANs for a Lossy Channel," *IEEE Commun. Lett.*, vol. 9, no. 2, Feb. 2005, pp. 100–02.

[7] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE JSAC*, vol. 18, no. 3, Mar. 2000, pp. 535–47.

[8] R. Nelson, Probability, *Stochastic Processes, and Queueing Theory*, Spinger-Verlag, 1995.

[9] X. Ling, Y. Cheng, J. W. Mark, and X. Shen, "A General Analytical Model for the IEEE 802.15.4 Contention Access Period," *Proc. IEEE WCNC*, Mar. 2007.

[10] L. Kleinrock and F. Tobagi, "Packet Switching in Radio Channels: Part I-Carrier Sense Multiple-Access Modes and Their Throughput-Delay Characteristics," *IEEE Trans. Commun.*, vol. 23, no. 12, Dec. 1975, pp. 1400–16.

[11] IEEE 802.15.4, "Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks (LR-WPANs)," 2003.

[12] X. Ling *et al.*, "A Novel Performance Model for Distributed Prioritized MAC Protocols," *Proc. IEEE GLOBECOM '07*, accepted.

Biographies

XINHUA LING (x2ling@bbcr.uwaterloo.ca) received a B.Eng. degree in radio engineering from Southeast University, China, in 1993, an M.Eng. degree in electrical engineering from the National University of Singapore in 2001, and a Ph.D. degree in electrical and computer engineering from the University of Waterloo, Canada, in 2007. Since May 2007 he has been a post-doctoral fellow at the University of Waterloo. His research interests are in the area of cellular, WLAN, WPAN, ad hoc networks, and their internetworking, including MAC protocol design and performance analysis, and cross-layer design for QoS support.

YU CHENG [S'01, M'04] (cheng@iit.edu) received B.E. and M.E. degrees in electrical rngineering from Tsinghua University, Beijing, China, in 1995 and 1998, respectively, and a Ph.D. degree in electrical and computer engineering from University of Waterloo, Canada, in 2003. From September 2004 to July 2006 he was a postdoctoral fellow in the Department of Electrical and Computer Engineering, University of Toronto, Canada. Since August 2006 he has been an assistant professor in the Department of Electrical and Computer Engineering, Illinois Institute of Technology, Chicago. His research interests include application-oriented networking, Internet performance analysis and resource allocation, network measurement, wireless networks, and wireless/wireline interworking. He received a Postdoctoral Fellowship Award from the Natural Sciences and Engineering Research Council of Canada (NSERC) in 2004.

JON W. MARK [M'62, SM'80, F'88, LF'03] (jwmark@bbcr.uwaterloo.ca) received a Ph.D. degree in electrical engineering from McMaster University, Canada in 1970. Upon graduation he joined the Department of Electrical Engineering (now Electrical and Computer Engineering) at the University of Waterloo, became a full professor in 1978, and served as Department Chairman from July 1984 to June 1990. In 1996, he established the Centre for Wireless Communications (CWC) at the University of Waterloo and has since been serving as the founding director. His current research interests are in wireless communications and wireless/wireline interworking, particularly in the areas of resource management, mobility management, and end-to-end information delivery with QoS provisioning. He is a co-author of the textbook *Wireless Communications and Networking* (Prentice-Hall, 2003). He has served as a member of a number of editorial boards, including *IEEE Transactions on Communications*, *ACM Wireless Networks*, and *Telecommunication Systems*. He was a member of the Inter-Society Steering Committee of *IEEE/ACM Transactions on Networking* from 1992 to 2003, and a member of the IEEE ComSoc Awards Committee during the period 1995–1998.

XUEMIN (SHERMAN) SHEN [M'97, SM'02] (xshen@bbcr.uwaterloo.ca) received a B.Sc. (1982) degree from Dalian Maritime University, China, and M.Sc. (1987) and Ph.D. (1990) degrees from Rutgers University, New Jersey, all in electrical engineering. He is a professor and University Research Chair, and associate chair for graduate studies, Department of Electrical and Computer Engineering, University of Waterloo, Canada. His research focuses on mobility and resource management in interconnected wireless/wired networks, UWB wireless communications systems, wireless security, and ad hoc and sensor networks. He is a co-author of three books, and has published more than 300 papers and book chapters in wireless communications and networks, control and filtering. He serves as the Technical Program Committee Chair for IEEE GLOBECOM '07, General Co-Chair for Chinacom '07 and QShine '06, the Founding Chair for IEEE Communications Society Technical Committee on P2P Communications and Networking. He also serves as a Founding Area Editor for *IEEE Transactions on Wireless Communications;* Editor-in-Chief for *Peer-to-Peer Networking and Application;* and Associate Editor for *IEEE Transactions on Vehicular Technology*, *KICS/IEEE Journal of Communications and Networks*, *Computer Networks*, and *ACM/Wireless Networks; and Wireless Communications and Mobile Computing* (Wiley). He has also served as Guest Editor for *IEEE JSAC*, *IEEE Wireless Communications*, and *IEEE Communications Magazine*. He received the Excellent Graduate Supervision Award in 2006, and the Outstanding Performance Award in 2004 from the University of Waterloo, the Premier's Research Excellence Award (PREA) in 2003 from the Province of Ontario, Canada, and the Distinguished Performance Award in 2002 from the Faculty of Engineering, University of Waterloo. He is a registered Professional Engineer of Ontario, Canada.