

Fairness Issues in Hybrid 802.11b/e Networks

Arvind Swaminathan and Jim Martin

{swamina,jmarty}@clemsun.edu

Clemson University, Clemson, SC 29634

Abstract—In this paper we evaluate the performance of a hybrid WLAN network which consists of a mixture of 802.11e nodes and 802.11b nodes. Using simulation, we show that the choice of the contention parameters severely affects the trade-off between the performance improvement at the 802.11e nodes and the performance degradation at the 802.11b nodes. An imperfect choice of contention parameters could lead to throughput starvation at the 802.11b nodes or to a severe degradation in the ability of the 802.11e nodes to meet QoS guarantees. We demonstrate that the best choice of contention parameters varies with the actual mix of 802.11e and 802.11b nodes and we propose an adaptive algorithm that addresses the issue.

I. INTRODUCTION

The protocols used in WLANs are defined under the IEEE 802.11 family of standards. These standards specify the air-interface between a wireless client and a base station or between two wireless clients in a WLAN. The air-interface includes both the modulation/coding used at the physical layer and the protocols used at the medium-access layer to allow the nodes in the WLAN to share the common wireless channel. An overview of the standards can be found in [1].

The current 802.11 standards are suitable only for carrying best-effort internet traffic. In order to support traffic with real-time constraints such as VoIP, the MAC layer must provide techniques to differentiate various classes of traffic, and there must be a method by which applications can specify their requirements and these resources be reserved at various points in the network. The 802.11e working group [7] was started in order to accomplish this goal of supporting time-sensitive traffic on WLANs.

The 802.11e standard is backward compatible with the 802.11b standard. It is likely that initial deployments of 802.11e will occur in hybrid networks that contain a mixture of 802.11e and 802.11b nodes. Our concern, and the motivation for the research we report in this paper, is the impact that the 802.11b and 802.11e nodes have on each other in hybrid WiFi networks. Through simulation we show that an imperfect choice of contention parameters could lead to throughput starvation at the 802.11b nodes or to a severe degradation in the ability of the 802.11e nodes to meet QoS guarantees.

This paper is organized as follows. First we provide an appropriate discussion on the 802.11b and 802.11e standards. We then analyze a hybrid network that contains both 802.11e nodes and 802.11b nodes. We then describe related work followed by concluding remarks and directions for future work.

II. BACKGROUND

A. 802.11b MAC Protocol

The 802.11b standard defines two coordination functions. The first is the Distributed Coordination function (DCF) and the second is the Point Coordination function (PCF). PCF is not widely deployed [3],[6] and hence we restrict our discussion to DCF.

This work was supported in part by the DoD Multidisciplinary University Research Initiative (MURI) program administered by the Office of Naval Research under Grant N00014-00-1-0565.

The DCF mode of operation is an adaptation of the CSMA/CA protocol. Consider the time at which the channel just becomes free as shown in figure 1. Each of the nodes that have a packet available waits for an interval denoted by Data Inter-Frame Space (DIFS). After this, the nodes wait for a random number of slots. The number of slots is obtained by picking a random number between $[1, CW+1]$, where CW is the size of the congestion window. The nodes wait for the number of slots remaining since the last time they picked a random number of slots. Figure 1 shows that node 1 has 4 slots left while node 2 has 6 slots left. The nodes then begin decreasing their contention window (i.e., the number of slots they have to defer).

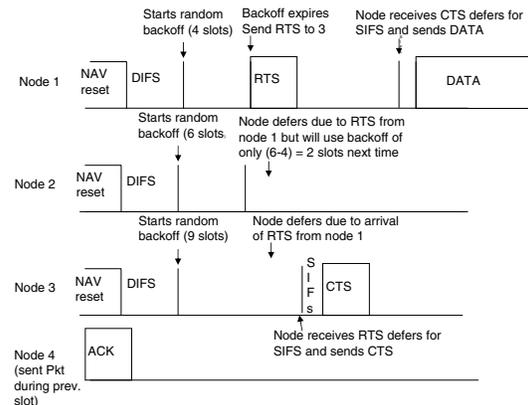


Fig. 1. Exchange of packets using the Distributed Coordination Function.

Once the contention window of a node goes to zero the node transmits an RTS to the destination (node 1 in the figure). The destination then responds with a CTS after an interval of time referred to as the Short Inter-Frame Space (SIFS). The nodes which overhear these control messages store the value of the contention window that they had counted down to. They also note the duration of the current data exchange (contained in the RTS and CTS messages) and defer transmissions until that point in time. The transmission concludes when the ACK sent from the receiver is received at the source of the RTS.

If the contention window of two different nodes goes to zero at the same time, both the nodes will transmit a RTS resulting in both packets being destroyed due to collision. In this case the contention window CW is multiplied by two and the nodes chose a new value for the contention window by choosing a random number over this expanded window.

B. Providing Quality of Service in 802.11 Networks

References [2] and [3] describe numerous techniques to improve the basic 802.11b standard to support real-time traffic. Methods that offer service differentiation can be classified into priority-based

methods and fair scheduling based methods. Priority-based methods differentiate traffic by using different values for the contention parameters for the different traffic classes. In the fair-scheduling based methods, differentiation is obtained by assigning different weights to different classes of traffic and calculating the wait times based on these weights. In both cases, the main contention parameters that are varied are the size of the initial contention window, the inter-frame spacing times and the parameters of the back-off algorithm.

The general consensus [2] is that fair-scheduling based schemes are advantageous as they can prevent starvation of low priority classes of traffic. However, they require substantial changes to the basic DCF algorithm. The priority based schemes require less change to the basic DCF and hence the recently proposed 802.11e standard [7] is based on priority-based methods.

C. 802.11e MAC Protocol

The 802.11e standard improves the MAC layer to support the performance requirements of multimedia and voice traffic. In order to provide QoS support, the 802.11e standard has two additional coordination functions. These are the Enhanced DCF (EDCF) and the Hybrid Coordination Function (HCF). The two newer modes of operation operate on traffic classified into one of 8 different priority levels (referred to as traffic classes or TCs).

In the EDCF mode of operation, MAC layer parameters are adapted based on the particular TC. The first parameter used to differentiate higher priority traffic is the interframe spacing interval (called the Arbitration Inter-frame spacing [AIFS] in 802.11e) assigned to the various TCs. As shown in figure 2, this parameter is made smaller for TCs with higher priority and thus we obtain a deterministic priority mechanism.

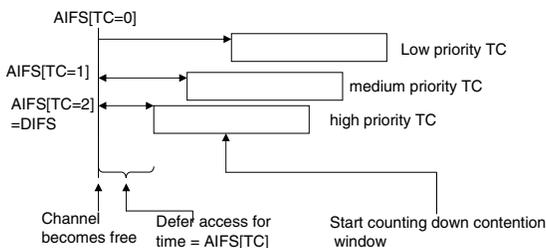


Fig. 2. Traffic Differentiation using AIFS parameter in EDCF.

After the AIFS, the contention period (CP) starts. Each TC within the station independently starts a back-off counter and contends for transmission opportunities (TxOp). The second parameter used to prioritize traffic is the initial minimum value of the back-off counter. This value is determined by the minimum collision window (CW_{min}) which is the number of slots a node backs-off. A smaller value of CW_{min} is assigned to traffic of higher priority. When a collision occurs, the value of the contention window CW is multiplied by a factor called the persistence factor (PF). A smaller PF indicates higher priority. Finally, the value of the contention window is not increased beyond a maximum value referred to as CW_{max} which also is smaller for higher priority traffic.

The differences between the DCF and the EDCF are illustrated in figure 3. There is only one queue at each node in the DCF mode. On the other hand, in the EDCF mode, the traffic is classified into eight different classes. Each class has its own set of parameters. These parameter values are chosen to provide the necessary amount

of differentiation among the traffic classes. The scheduler is used to resolve a “virtual collision”, which occurs when the back-off timer for two different classes of traffic within the same node expire at the same time. The scheduler assigns the TxOp to the higher priority traffic class.

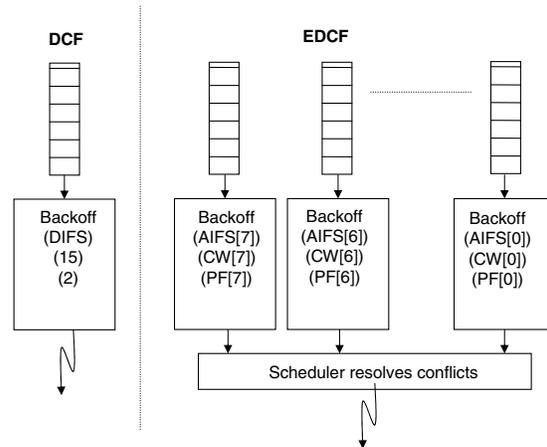


Fig. 3. DCF mode of operation versus EDCF mode of operation.

III. PERFORMANCE EVALUATION

In a hybrid WiFi network, 802.11b nodes will continue to use DCF mode while 802.11e nodes will use EDCF mode. One way to model a hybrid network is by designating 802.11b nodes as 802.11e nodes configured to use the priority setting of the lowest priority class.

Several issues arise in a hybrid network. The first is that since 802.11e nodes employ a priority-based mechanism to differentiate traffic classes, it is possible that the parameters of the high priority traffic are set in a manner that leads to throughput starvation of the lower priority traffic class (and hence at the 802.11b nodes). The second is that if the parameters are not tuned properly, the higher priority traffic in the 802.11e nodes might start to see a decrease in performance as the amount of lower priority traffic (or as the number of 802.11b nodes) increases. We conjecture that by varying the values of the contention parameters one should be able to obtain a good trade-off between the penalty suffered by high priority traffic in the 802.11e nodes and the throughput degradation at the 802.11b nodes. We use delay as the metric to quantify the impact of varying the contention parameters.

A. Methodology

We modified the commercially available OPNET 802.11b model to support hybrid networks. The RTS/CTS model is not used and instead the hidden terminal problem is addressed by issuing a busy signal to all the nodes once a data transmission starts at a particular node. This simplifying assumption was made because it closely approximates an enterprise WLAN in which the nodes are within range of each other. We also assume that if two nodes transmit at the same time, neither packet is received at the access point (AP), which is justified by the low processing gain employed in WLANs. Our analysis considers a home network environment involving a small number of clients generating best effort Internet traffic along with a controlled amount of VoIP or multimedia traffic that requires higher priority access. We also assume that all the nodes have equally good channels to the AP.

TABLE I

PARAMETER VALUES USED TO STUDY THE EFFECTIVENESS OF VARIOUS CONTENTION PARAMETERS.

Differentiating Parameter	CW_{min}		PF		AIFS		CW_{max}	
	Priority		Priority		Priority		Priority	
	Low	High	Low	High	Low	High	Low	High
None	6	6	2	2	2	2	64	64
AIFS	6	6	2	2	4	1	64	64
CW_{min}	6	2	2	2	1	1	64	64
PF	6	2	2	1.4	1	1	64	64
CW_{max}	6	6	2	2	1	1	64	16

TABLE II

PARAMETER VALUES USED IN THE STUDY TO OBTAIN THE “BEST” CONTENTION PARAMETERS FOR A HOME NETWORK.

Case	CW_{min}		PF		AIFS		CW_{max}	
	Priority		Priority		Priority		Priority	
	Low	High	Low	High	Low	High	Low	High
1	6	6	2	2	2	2	64	64
4	8	2	2	2	1	1	64	64
7	8	2	2	2	4	1	64	8
9	6	6	2	2	4	1	64	64
13	6	2	2	2	5	1	16	16

In the simulation model nodes employ a constant-bit rate source that generates packets at a rate of 256 kbps. The channel bit rate is taken to be 1Mbps. The packets are queued at the network layer and are transmitted in order based on the availability of the channel. The simulations continue until 50,000 packets are generated at each node. Five runs are simulated, each with a different starting seed. For each packet generated at a node, the time interval between the insertion of the packet in the queue until the arrival of the ACK for the packet is measured. These values are averaged to give the mean delay of the packets at that node for that particular run. The mean delay over all the runs is then obtained by averaging the delay obtained with the five different starting seeds. The contention parameters for the p th traffic class are: initial contention window ($CW_{min}[p]$), maximum size of the contention window ($CW_{max}[p]$), interframe spacing ($AIFS[p]$), and the factor used to increase contention window when collisions occur ($PF[p]$).

B. Effect of Contention Parameters

To quantify the effectiveness that $AIFS$, CW_{min} , CW_{max} , and PF have on achieving traffic differentiation, we first consider the performance of a scenario that employs a single 802.11e node along with a smaller number of competing 802.11b nodes. The number of competing 802.11b nodes is increased from 1 to 9. We simulated the network for the parameter settings shown in table I. The values in the high priority column refer to the contention parameters of the 802.11e node while the values under the low priority column refer to the contention parameters employed by the 802.11b nodes.

The mean value of the end-to-end delay associated with successful packets sent by the 802.11e node is shown in figure 4. The curve labeled “No differentiation” is the baseline case as it represents all nodes configured identically using reasonable 802.11b contention parameters. As the number of competing 802.11b nodes increases, we see the delay grow exponentially and then linearly once the network becomes saturated. As we can see from figure 4, the effect of using PF and CW_{max} to differentiate the 802.11e node traffic from the 802.11b node traffic leads to a noticeable improvement in performance of almost a factor of two. But still, as the number of competing 802.11b nodes is increased, the mean delay worsens by almost an order of magnitude. AIFS is the most effective parameter. Figure 4 demonstrates that assigning a higher value of AIFS to the low priority traffic leads to an order of magnitude improvement in performance when compared to the case when no differentiation is employed. The mean delay obtained with AIFS based differentiation is lower than the mean delay obtained with CW_{min} (the next most effective parameter) based differentiation by a factor of four. The 802.11e node delay increases only by a factor of two as the number of competing 802.11b nodes is increased.

Figure 5 shows the CDF of the mean delay for the case in which there is one 802.11e node and nine 802.11b nodes. As we can see

the CDF goes to 1 most quickly for the AIFS based differentiation and least quickly for the case when there is no differentiation. With the AIFS based scheme, 90% of the packets experience a delay of 0.02s or less while 100% packets experience a delay of 0.1s or less. Without differentiation, 90% of the packets experience a delay of 0.75s or less. Additionally, we found that the traffic differentiation does not affect the delay performance of the 802.11b nodes appreciably. (We have not shown this figure due to space constraints and also because we will see a similar but much stronger result in the next section.)

C. Determining the “best” contention parameters

Previous results suggest that an 802.11e session is highly sensitive to the choice of contention parameters as the number competing 802.11b nodes is varied. Motivated by this, we examined the effect of various combinations of 802.11e contention parameters in an effort to improve the delay performance of the 802.11e node. The best performing cases are listed in table II (case 1 has the same parameters as the no differentiation case in the previous section). The corresponding delay performance and CDF plots are shown in figure 6-8.

Figure 6 illustrates that the contention parameters used in cases 7 and 13 are extremely successful at isolating the 802.11e node traffic from increasing levels of 802.11b traffic. The increase in mean delay as the increase in the number of competing 802.11b nodes is almost indiscernible. We conclude that it is possible to select contention parameters such that resulting traffic differentiation makes the 802.11e node performance almost immune to the number of competing 802.11b nodes.

The curves in figure 7 show the performance of the 802.11b nodes for the same set of contention parameters. With the parameters in case 13 the performance of the 802.11b nodes is almost identical to the baseline case, suggesting that the performance improvement at the 802.11e node comes at virtually no cost to the performance of the 802.11b nodes. Finally, the curves in figure 8 show the CDF of the delay experienced by the successful packets for the various cases. As seen from the CDF curve, if the parameters in case 13 are employed, 100% of the packet delays are below 0.01 seconds. Also, the delay values have an average value of 0.0075s and 95% of the delay values falls in the range [0.006, 0.0095].

The results in this subsection suggest that for a home network with one 802.11e node competing with several 802.11b nodes, an optimal set of contention parameters exist such that the delay experienced by the 802.11e node has a low average value and minimal jitter. Further, this behavior is invariant to a reasonable number of competing 802.11b nodes and comes at virtually no cost to the 802.11b node performance. This insight can be extended to an all 802.11e network. In this environment it is possible to choose contention parameters to achieve ideal performance when a single high priority stream competes with multiple low priority streams.

D. Effect of percentage of 802.11e nodes and 802.11b nodes

Although the results in the previous sub-section are extremely promising, in a practical network the ratio of the number of 802.11b nodes to the number of 802.11e nodes might vary in a manner different from the cases considered previously. We next consider the effects of the various contention parameters on performance in a network in which the ratio of the number of 802.11b nodes to that of the 802.11e nodes is maintained at 1. We use the contention parameters given in the table in figure I. Figure 9 shows the number of successful packets transmitted by the 802.11e nodes while figure 10 shows the number of successful packets transmitted by the 802.11b nodes.

Figure 9 demonstrates that the AIFS based differentiation still leads to substantially improved performance when the performance of 802.11e nodes is considered. The number of successfully transmitted packets almost doubles with AIFS based differentiation. However, figure 10 suggests that the level of performance achieved by the 802.11b nodes becomes unacceptable with the AIFS based differentiation scheme. The 802.11b nodes suffer from throughput starvation. The number of packets successfully transmitted with the other differentiation schemes is comparable to that obtained with no differentiation. This suggests that to ensure fairness one must adapt the contention parameters being varied based on the mix of the 802.11b and 802.11e nodes.

IV. RELATED WORK

References [2] [3] [6] describe the limitations of the 802.11b standard in supporting real-time applications with stringent delay and jitter requirements. They also describe techniques that have been proposed to provide guaranteed QoS to applications in 802.11 networks, and describe the 802.11e standard [7] in detail. Reference [4] investigates the performance obtained with the new standard. However, no attempt is made to optimize any of the parameters. Reference [5] considers a scenario in which a high priority traffic stream competes with a low priority traffic stream. The values of two contention parameters are fine tuned to improve the performance of the high priority stream. However, the fine tuning only considers the effect on the performance of the high priority stream and not the penalty suffered by the low priority stream. Moreover, they do not consider the fine tuning with multiple traffic streams. To the best of our knowledge ours is the first work that analyzes the effect of varying the different contention parameters on the trade-off between the performance improvement at the high priority streams and the performance penalty suffered by the low priority streams.

V. CONCLUSIONS AND FUTURE WORK

We have developed an OPNET simulation model that enables us to compare the effect of the various parameters used to differentiate high priority traffic from the low priority traffic in 802.11e networks. We have shown that there is a tradeoff between the level of isolation achieved by an 802.11e node and the impact of this on competing 802.11b nodes. The results demonstrate that due to the deterministic nature of traffic differentiation provided by AIFS based differentiation, it results in the best delay performance for high priority traffic in the presence of many competing low priority streams. On the other hand, PF based differentiation leads to greater fairness in the trade-off between the performance improvement for the high priority traffic and the performance penalty to the low priority traffic.

The results also demonstrate that for a network in which a single 802.11e node competes with multiple 802.11b nodes, it is possible

to chose contention parameters such that the traffic from 802.11e node experiences low delay with almost zero jitter, while the delay performance of the traffic from 802.11b nodes is comparable to the case when there is no traffic differentiation. However, this same set of contention parameters lead to throughput starvation of the 802.11b nodes when the hybrid network consists of a greater percentage of 802.11e nodes.

We are currently addressing this issue by tabulating [8] the performance with different contention parameters for networks with an arbitrary mix of low and high priority traffic streams. This table can be used in two ways [8]. In the first technique, the contention parameters listed in the table that provide sufficient fairness in a network with a single low priority stream and multiple high priority streams are used always. This will guarantee a certain degree of fairness under all scenarios (we are designing for the worst-case). The second technique is an adaptive technique in which the AP first estimates the traffic mix in the network. Following this, it uses the above lookup table to choose the set of contention parameters to be used by the nodes and issues a control command with these parameters. The parameters are chosen based on the estimated traffic mix and a performance measure to be optimized. The measure in general can be a combination of fairness and QoS requirements of the high priority streams. We also plan to investigate the effect of having more than two priority levels.

REFERENCES

- [1] M. Davis, "Overview of 802.11 Family of WLAN standards," available online at http://www.electronics.dit.ie/staff/mdavis/IEE_Talk_040304.pdf.
- [2] Hua Zhu et. al., "A survey of Quality of Service in IEEE 802.11 Networks," *IEEE Wireless Comm. Mag.*, August 2004
- [3] G. Chesson, "An Overview of Quality of Service issues in Wireless Networks," *Lecture available online at <http://stanfordonline.stanford.edu/courses/ee380/010207-ee380-100.aspx>*
- [4] Stefan Mangold et. al., "IEEE 802.11e Wireless LAN for Quality of Service," *Proceedings of the European Wireless 2002 conference*
- [5] Priyank Garg et. al. "Using IEEE 802.11e MAC for QoS over Wireless," *Proceedings of IPCCC 2003*
- [6] Yang Xiao, "IEEE 802.11e: QoS Provisioning at the MAC layer," *IEEE Wireless Comm. Mag.*, June 2004
- [7] IEEE 802.11 WG, "Draft Supplement to Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: MAC Enhancements for Quality of Service," IEEE Std 802.11e/D3.3.2, Nov. 2002
- [8] A. Swaminathan, "Investigating fairness Issues in Hybrid 802.11b/e Networks," CPSC 811 Final project report available at <http://www.clemson.edu/~swamina/881Report.pdf>

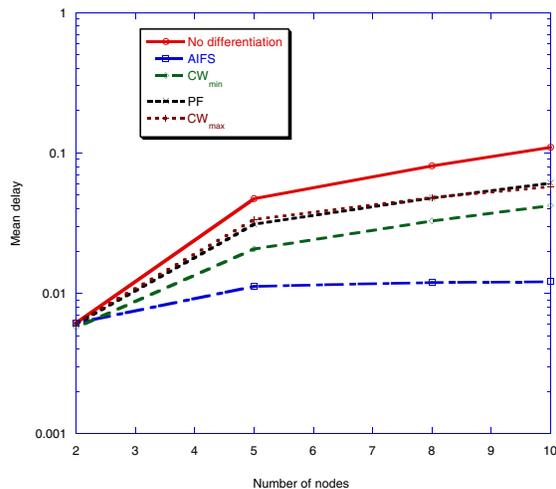


Fig. 4. Effectiveness of different contention parameters in providing traffic differentiation.

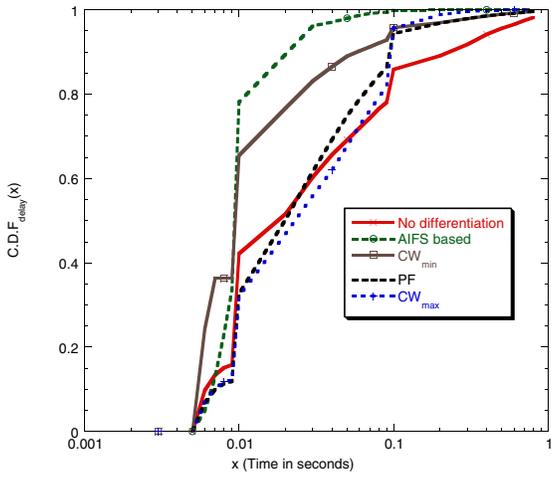


Fig. 5. CDF of delay at the 802.11e node for traffic differentiation based on the different contention parameters.

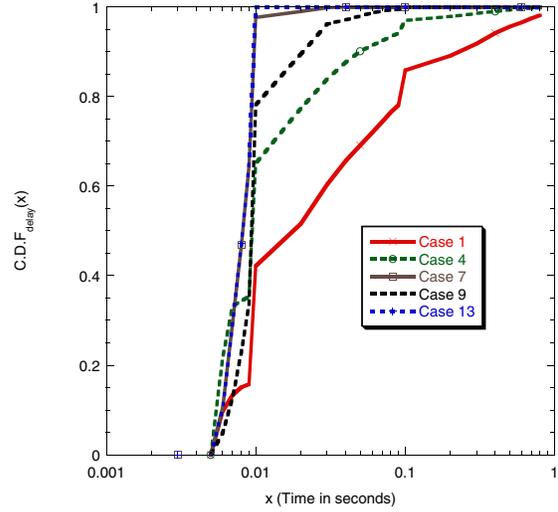


Fig. 8. CDF of delay at the 802.11e node with the different contention parameters.

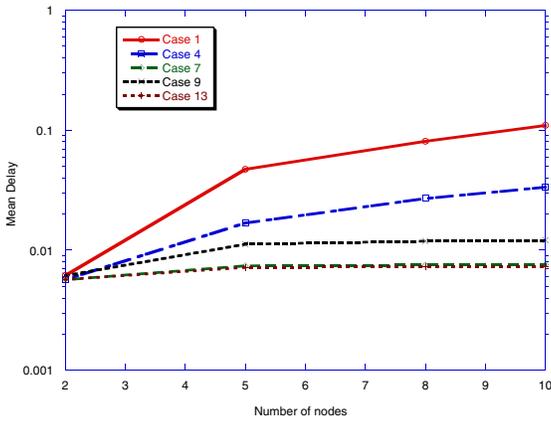


Fig. 6. Optimizing the performance of the 802.11e node by a proper choice of the contention parameters.

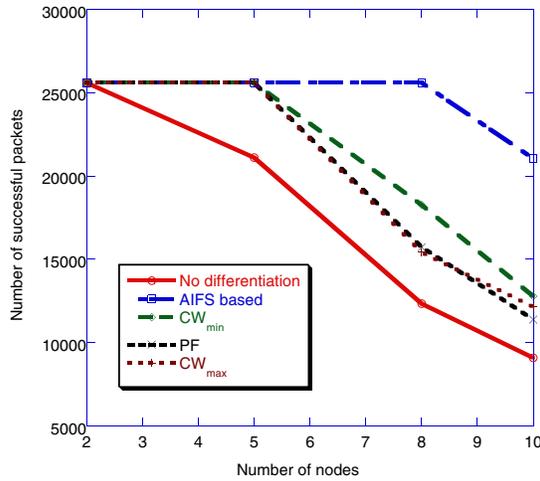


Fig. 9. Number of successful packets transmitted by the 802.11e nodes for various choices of the contention parameters.

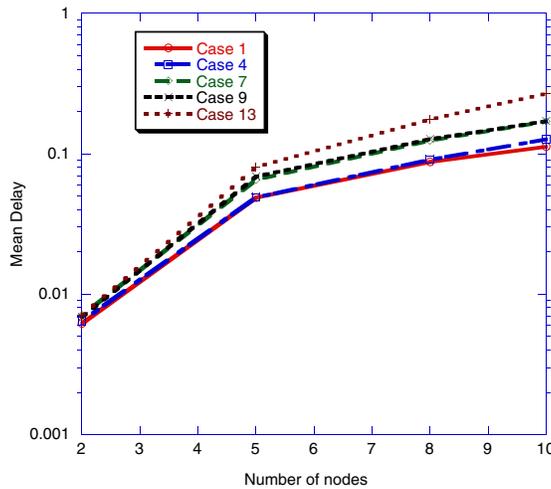


Fig. 7. Performance of the 802.11b nodes for the various choices of the contention parameters.

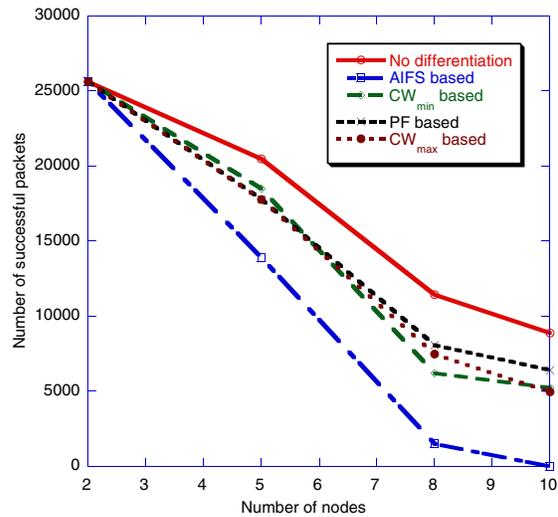


Fig. 10. Number of successful packets transmitted by the 802.11b nodes for various choices of the contention parameters.