

# A Station-Based Adaptation Algorithm to Improve Robustness of IEEE 802.11

Jean Cao      Yujing Wu      Carey Williamson  
Department of Computer Science, University of Calgary  
Email: {caox,ywu,carey}@cpsc.ucalgary.ca

## Abstract

*This paper studies a Bad Apple phenomenon caused by Head-of-Line blocking in IEEE 802.11 networks. This problem can adversely affect the performance of wireless multimedia streaming. We summarize the characteristics of the Bad Apple problem, and then propose a Station-Based Adaptation (SBA) algorithm to solve it. SBA maintains per-station information, operating at the MAC layer as an extension to 802.11 DCF. SBA deactivates flows destined to temporarily unreachable stations, and later reactivates flow transmission when the stations are available. Simulation results show that SBA can effectively improve the performance and robustness of the 802.11 system when stations experience unpredictable temporary failures that are several seconds in duration.*

## 1 Introduction

IEEE 802.11 wireless local area networks (WLANs) [9] have been widely deployed in hot spots, homes, universities, and business offices. As a popular “last mile” technology, 802.11 has been increasingly used by multimedia applications, such as video conferencing, media streaming, and Internet telephony. As multimedia traffic volumes increase in WLANs [7], high rate continuous packet flows pose new challenges to current 802.11 networks.

Consider a scenario in which several clients are concurrently receiving video streams from a server in an 802.11 ad hoc network. If the channel conditions for one client become too poor to receive frames (e.g., the client moves too far away), then all other clients can lose their playback [6]. Their playback resumes only when the “bad” station reconnects. In our earlier experiments [3, 6], we observed this phenomenon and called it the Bad Apple (BA) phenomenon after the saying “One bad apple spoils the batch”.

As discussed in [6], the Bad Apple phenomenon is caused by Head-of-Line (HOL) blocking at the server’s FIFO queue in the data link layer. Even when a station is temporarily unreachable, the server still keeps retransmitting frames to it. In the IEEE 802.11 MAC protocol,

a frame undergoes up to MAX\_RETRY (maximum retry limit; 802.11 suggests 7 for short frames, and 4 for long frames, whereas on Cisco Aironet 350 NIC, 16 was used) retransmissions, with increased waiting times (i.e., the binary exponential backoff (BEB) algorithm) between each attempt. During such a retransmission episode, frames to other stations are not transmitted or delivered. This may cause excessive queuing delay and frame loss at the server, especially for multimedia flows [3].

Temporary transmission failures in 802.11 WLANs are inevitable due to user mobility and random interference from the surrounding environment. These problems can seriously degrade the overall system throughput, and may violate the Quality of Service (QoS) requirements for other flows in the network. Greater robustness is needed for QoS support in IEEE 802.11 WLANs.

Our paper makes two main contributions. First, we carefully characterize the Bad Apple phenomenon, identifying the conditions under which it occurs, and the effects on system performance. Second, we propose and evaluate a solution to the Bad Apple problem.

The first part of the paper studies the Bad Apple problem with extensive ns-2 simulations, and shows that the preconditions of the BA problem are the following:

1. Multiple stations share a single FIFO transmission queue at the data link layer.
2. The MAC-layer protocol uses multiple retransmissions.
3. High-rate UDP traffic is sent to a station that is unreachable for several seconds.

In the second part of the paper, we propose a Station-Based Adaptation (SBA) algorithm to solve this problem. Per-station information is used to control the transmission to each station based on its channel availability. If a station is unreachable, flows destined to this station are disabled, and resumed when the station is reachable again.

The rest of the paper is organized as follows. Section 2 reviews related work. Section 3 illustrates the impacts of the BA phenomenon. The SBA algorithm is presented in Section 4, and its performance is evaluated using simulation in Section 5. Section 6 concludes the paper.

## 2 Related Work

The BA problem is triggered by HOL blocking, which typically occurs in FIFO queues. The HOL problem has been studied for many years, primarily in wired networks.

The HOL problem also manifests itself in wireless networks. In 1996, Bhagwat *et al.* [5] proposed a link-layer channel-state-dependent packet (CSDP) scheduling scheme for WLANs. The WLAN in their study is not the same as the current 802.11 WLANs, but the fundamental features are the same.

CSDP operates multiple per-station queues at the Access Point (AP) in a round robin manner. Upon a packet transmission failure, the queue corresponding to the destination is not allowed to send packets in the next burst duration. CSDP relieves the impact of HOL blocking when the channel failure occurs at short timescales, i.e., hundreds of milliseconds. However, it cannot effectively handle longer term disconnections (e.g., several seconds), because it is difficult to determine a suitable deferring time in practice. Moreover, after each fixed length deferring duration, the frame destined to the bad station is sent to the MAC layer for retry, which takes a long time due to 802.11's BEB algorithm and retransmission policy. Bandwidth is wasted by delivering significantly delayed multimedia frames.

Recently, Jiang and Liew [10] improve CSDP by applying an adaptive round robin (ARR) scheduler. The first frame in the queue corresponding to the bad station is deferred for transmission for a number of “rounds” instead of a fixed duration used by CSDP. Every transmission failure (up to MAX\_RETRY retransmissions) doubles the number of “rounds” that the queue is skipped before the next retry. Also, they modify the 802.11 MAC by using a fixed contention window size  $W$  instead of BEB that is used in 802.11. The intention is to shift control from the MAC layer to the link layer, so that the unlucky frame defers its transmission multiplicatively.

ARR still has the aforesaid limitations. The bandwidth is wasted by retransmitting frames destined to the unreachable destination, and delivering outdated frames. Moreover, this algorithm causes unfairness to the uplink flows (transmission from a mobile station to the AP). When a collision happens, a frame at the AP defers only time  $R$  (a random number between  $[0, W-1]$ ), whereas the mobile station with the standard 802.11 has to defer based on BEB.

Xia and Hamdi [12] have proposed a cross-layer system called WFS-ARC that uses station-based information for LLC layer scheduling decisions. WFS-ARC is a promising solution to deal with the HOL blocking problem in a multi-rate WLAN. However, it does not consider channel disconnection, in which case retransmitting at lower rates worsens the HOL blocking problem.

Previous work focuses on HOL blocking caused by poor




Figure 1. Simulation Setup for BA Problem

channel conditions occurring at short time scales (e.g., milliseconds). In this paper, we study the BA phenomenon triggered by HOL blocking over longer time scales. A Station-Based Adaptation (SBA) algorithm is proposed to solve the BA problem.

## 3 Bad Apple Effects in 802.11 Networks

In this section, we explore the impacts of the BA problem. We define a *Bad* flow to be a flow destined to a station that is temporarily disconnected from the network. A *Victim* flow is a flow that shares the same transmission queue with the Bad flow. A *Background* flow is a flow that exists elsewhere in the same network, but does not share a transmission queue with the Bad flow.

We use ns-2 simulation to study the impacts of the Bad flow on the Victim flow, the Background flow, and the overall system performance.

Figure 1 shows the network topology used for the simulations. In infrastructure-based 802.11 networks, the Background flows could be uplink flows from mobile stations to the AP. A flow can be a UDP flow or a TCP flow. We use UDP to represent a constant bit rate video streaming flow.

In our simulation setup, F1 is the Bad flow from station S to Station C1. Station C1 is unreachable for 10 seconds (from time 30 to 40 seconds) during the 100-second simulation. F2 is the Victim flow from station S to C2, and F3 is the Background flow from station SX to CX.

The other simulation parameters are set as follows. MAX\_RETRY is set to 7 for short frames, and 4 for long frames, as suggested in the 802.11 standard. The queue limit is 200 frames. The channel capacity is set to 1 Mbps, which provides about 750 kbps of usable bandwidth (due to the use of RTS/CTS). Because the main purpose is to evaluate the impacts of the BA problem, we do not overload the channel in most of our tests, though in some tests, the TCP flows saturate the channel.

The performance metrics are goodput, instantaneous queue length, frame drop ratio, and the average queueing delay. The frame drops are counted from simulation time 30 to 50 seconds. The average queueing delay is the average in-queue frame delay from time 30 to 50 seconds.

### 3.1 Effects of UDP Bad Apple Flow on Victim Flows

Figure 2 shows how the Bad Apple flow affects the Victim flow. The left column is for a UDP Victim flow in a network with a medium aggregate load of 400 kbps. The middle column shows a UDP Victim flow at a higher aggregate load of 600 kbps. The rightmost column shows the case in which the Victim flow is a TCP flow.

Figure 2(a) shows that the goodput of the Victim flow decreases during the Bad Apple period, and the queue shared by the Bad flow and the Victim flow fills. However, at medium load, no frame drops occur, since the Bad Apple flow resumes before the queue capacity (200 frames) is exceeded.

In a highly loaded network, as shown in Figure 2(b), the goodput of the Victim flow decreases dramatically when the Bad flow’s destination is unreachable. The queue shared by the Bad flow and the Victim flow fills soon after the disconnection, and the queue drops frames. These results match the experimental results in [6].

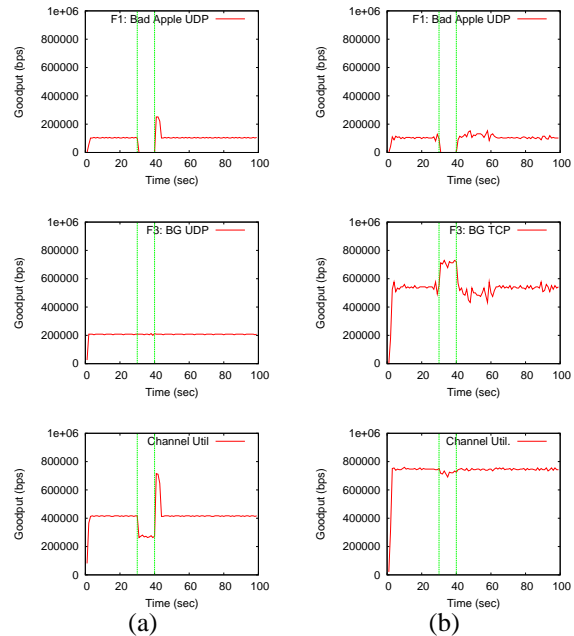
The results in Table 1 show that, in the presence of the BA problem, a Victim UDP flow could experience a high loss rate (around 24% in the test in Figure 2(b)) and long queueing delay (around 1.6 seconds for each frame). When the traffic load is lighter, as in the test in Figure 2(a), the HOL blocking does not cause frame loss. The user might experience a brief interruption of services, but the BA phenomenon does not occur.

When the Victim flow is a TCP flow (Figure 2(c)), the effects are slightly different. The goodput still degrades, but the frame drop ratio and the queueing delay are lower (see Table 1). The reason for the difference is TCP’s congestion control algorithm. Soon after the disconnection, the queue fills, and there are frame drops. These drops trigger TCP’s congestion control mechanism, which reduces the TCP flow rate. From that point on, the queue is dominated by frames from the Bad flow. Once the BA flow reconnects, the queue dissipates, thus allowing more successful transmissions for the Victim flow. Table 1 shows that the Victim flow experiences a low proportion of the frame drops.

### 3.2 Effects of UDP Bad Apple Flow on Background Flows

This section shows how the Bad flow affects the Background flow and overall system performance.

In Figure 3, the left column shows that if the Background flow has a constant bit rate, the BA problem does not affect the Background flow. The total channel utilization decreases due to the wasted channel for stations involved in the BA problem (i.e., S, C1, and C2). However, if the Background flow is a TCP flow, as in Figure 3(b), then the un-



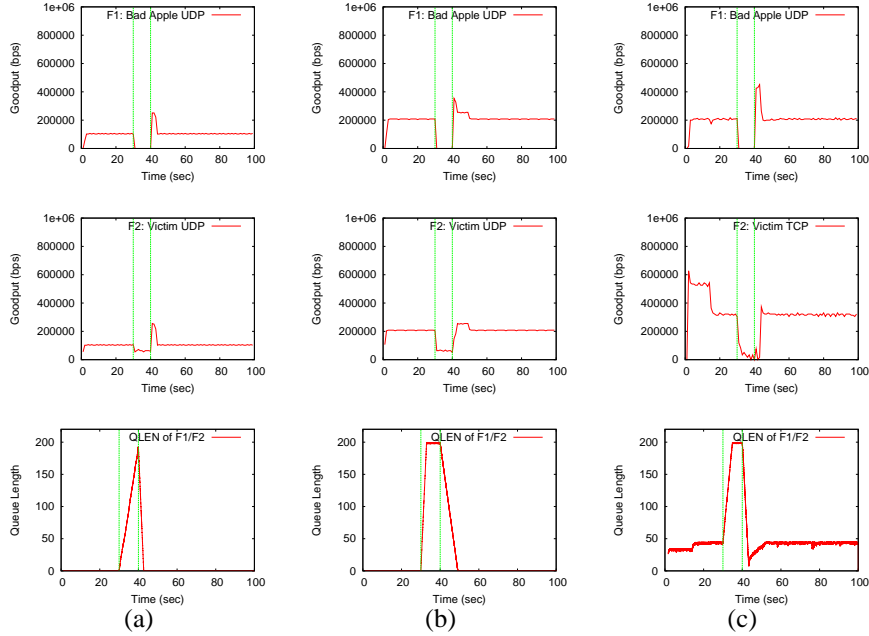
**Figure 3. Effects of UDP Bad Apple Flow F1 on Background Flow F3 (F1,F2: UDP 100 kbps): (a) F3: UDP 200 kbps; (b) F3: TCP**

used channel is available for the Background flow, and high channel utilization can be achieved.

Other tests (not shown here) consider multiple background flows. Multiple UDP Background flows can maintain their goodput during the disconnection. Similarly, multiple TCP Background flows share the extra bandwidth available, though not always in a fair manner.

### 3.3 Effects of TCP Bad Flow

A TCP flow should not cause the BA Phenomenon, due to the TCP congestion control. We demonstrate this in Figure 4. The left column shows that the TCP BA flow has minimal impact on the UDP Victim flow, the Background flow, and the overall system performance. The right column shows a case when the Background flow is also TCP. In this case, there is some degradation in goodput for the Victim flow during the disconnection. However, looking closer, we see that the goodput increases before the disconnection period ends. The reason is that when the Bad TCP flow experiences disconnection, it lowers its rate. The Victim flow experiences poor performance briefly, but is not blocked long. Therefore, this does not result in the BA problem.



**Figure 2. Effects of UDP Bad Apple Flow F1 on Victim Flow F2: (a) F1, F2: UDP 100 kbps; F3: UDP 200 kbps; (b) F1, F2, F3: UDP 200 kbps; (c) F1, F3: UDP 200 kbps; F2: TCP**

**Table 1. Frame Drop Ratio and Queuing Delay for Victim Flow**

Network Traffic Configuration (F1: Bad Apple; F2: Victim; F3: Background)	Frame Drop Ratio			Queue Delay (s)	
	Total	F1	F2	F1	F2
F1, F2: UDP 100 kbps; F3: UDP 200 kbps (Figure 2(a))	0.0%	0.0%	0.0%	1.24	1.25
F1, F2, F3: UDP 200 kbps (Figure 2(b))	<b>24.7%</b>	41.2%	58.8%	1.91	<b>1.61</b>
F1, F3: UDP 100 kbps; F2: TCP (Figure 2(c))	<b>9.9%</b>	97.3%	<b>2.7%</b>	2.27	<b>0.56</b>

## 4 Station-Based Adaptation Algorithm

As mentioned earlier, the Bad Apple problem is directly related to HOL blocking in longer time scale. We propose a Station-Based Adaptation (SBA) algorithm to solve this problem.

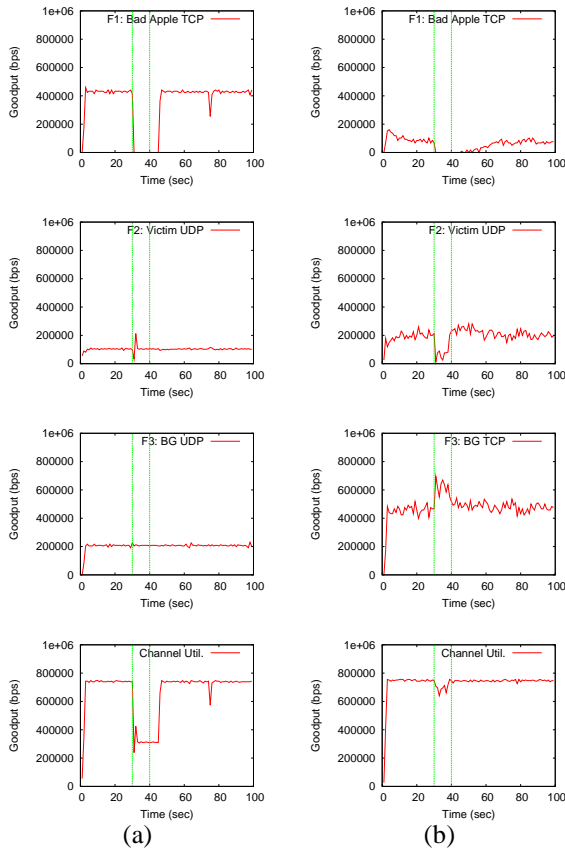
The basic idea of SBA is to monitor flows destined to each station. If a frame transmission failure occurs, SBA considers that the channel to this destination is getting worse. Consecutive failures to the same destination indicate longer term disconnection, and flows destined to it are deactivated. Later on, when the destination reconnects, the flows resume.

SBA uses a two-level adaptation scheme. In the first level, when the channel seems to be getting worse, the retry limit is reduced multiplicatively. This feature is meant to reduce the HOL effects caused by repeated retransmissions. When the channel seems to be getting better, the retry limit is increased multiplicatively. In the second level of con-

trol, when disconnection is detected, frames from the bad flow are transmitted only with a low probability. This limits the number of frames transmitted to the disconnected station. With a low but non-zero transmission probability, the flow can resume transmission when connectivity is re-established.

The main advantage of SBA is its simplicity. SBA extends the 802.11 MAC layer by using MAC information (such as destination ID), and operates in the MAC layer only. No cross-layer information is required. Moreover, SBA does not change the fairness properties of 802.11.

Another advantage of SBA is its low cost. It requires three per-station variables (i.e.,  $TX\_Prob$ ,  $Retry\_Limit$  ( $RL$ ), and  $T_{last\_prob}$ ), and three system-wide thresholds (i.e.,  $MIN\_TX\_PROB$ ,  $MIN\_RETRY$ , and  $TX\_PROB\_AGING$ ). The cost of storing such information in each wireless NIC is minor. The cost of monitoring the per-station variables is low, too. The parameters are updated upon certain IEEE 802.11 timeouts (indicating a transmission failure) or the receipt of



**Figure 4. Effects of TCP Bad Apple Flow F1: (a) F1: TCP; F2: UDP 100 kbps; F3: UDP 200 kbps; (b) F1,F3: TCP; F2: UDP 200 kbps**

ACKs (indicating a transmission success). No new timers are required.

Figure 5 illustrates how SBA reacts to channel disconnection. When a frame loss happens, SBA decreases the retry limit RL from MAX\_RETRY. The RL decreases if there are consecutive losses to the same destination, until it reaches MIN\_RETRY, which is set to 1 by default. The RL stays at 1 (no retry) as long as the destination is still disconnected. After a successful ACK is received, RL doubles after every ACK received until it reaches MAX\_RETRY. In the meantime, after the third frame loss, the TX\_Prob decreases multiplicatively from 1.0 down to MIN\_TX\_PROB. TX\_Prob is reset to 1.0 upon the receipt of the first successful ACK. TX\_Prob is also doubled after every TX\_PROB\_AGING interval.

Table 2 and Table 3 list the SBA parameters and their default values. Their roles are explained as follows:

**TX\_Prob and MIN\_TX\_PROB:** TX\_Prob controls the probability that a frame is sent to the physical channel for

**Table 3. System-wide Parameters in SBA**

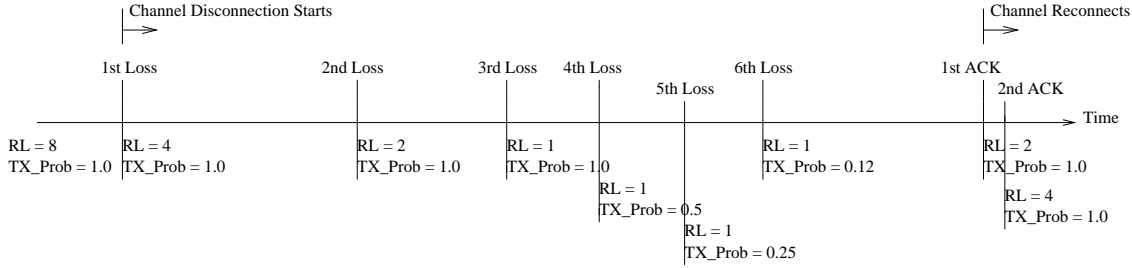
Parameters	Role	Default Value
MIN_TX_PROB	Lower bound for TX_Prob.	0.06
MIN_RETRY	Lower bound for retry limit.	1
TX_PROB_AGING	Aging period for resuming a flow.	30 sec.

transmission. To achieve the best performance, quickly turning off a bad flow is important. The effectiveness depends on how quickly a disconnected station can be detected, and how quickly TX\_Prob is decreased. Detecting a disconnected station is difficult in IEEE 802.11 (DCF) networks since there is no central controller to monitor the joining/leaving of stations. SBA detects a disconnected station based on repeated transmission failures. After three transmission failures, SBA considers the destination unreachable. TX\_Prob is decreased multiplicatively by a factor of 2 to disable the bad flow quickly.

MIN\_TX\_PROB is a system-wide threshold specifying the minimum legal value for TX\_Prob (i.e., the probability that a frame of the Bad flow is transmitted). There are two opposing considerations when setting its value. On one hand, a low MIN\_TX\_PROB saves channel bandwidth (e.g., a probability of 0 completely disables the bad flow). On the other hand, with MIN\_TX\_PROB = 0, there is no way to resume a bad flow when the disconnected station returns. A larger MIN\_TX\_PROB (e.g., 1.0) would resume the connection quicker. Currently, 0.06 is the default value in SBA, since it performs well (see Section 5.2.2).

**RL and MIN\_RETRY:** RL determines the retry limit. Frame losses due to short term disconnection can be recovered by retransmission, but for longer term disconnection, retransmissions can exacerbate HOL blocking and cause the Bad Apple problem. Reducing the retry limit can effectively reduce the Bad Apple effects [6]. In SBA, a multiplicative decrease/increase scheme (using a factor of 2) is used for changing RL. The RL value can be set as low as MIN\_RETRY, a system-wide threshold.

**$T_{last\_prob}$  and TX\_PROB\_AGING:**  $T_{last\_prob}$  remembers the last time that TX\_Prob was changed. This variable works in conjunction with TX\_PROB\_AGING. In particular, TX\_Prob is doubled every TX\_PROB\_AGING period since the last time TX\_Prob was changed ( $T_{last\_prob}$ ). Basically, the TX\_Prob value does not stay low for too long even if the traffic rate is slow, allowing the flow to be resumed sooner. This aging mechanism is especially useful for TCP Bad flows, since TCP congestion control reduces the transmission rate in the presence of frame loss. It could take the flow a long time to resume with a low



**Figure 5. Reactions of RL and TX\_Prob to Channel Disconnection in SBA Algorithm**

**Table 2. Per-Station Variables Used in SBA**

Parameters	Value Range	Init. Value	Update Conditions
TX_Prob	[MIN_TX_PROB, 1.0]	1.0	Halve after 3 TX failures. Reset to 1.0 after a successful TX.
Retry_Limit (RL)	[MIN_RETRY, MAX_RETRY]	1.0	Halve after a TX failure. Double after a successful TX.
$T_{last\_prob}$	[0, +∞)	0.0	Double TX_Prob every TX_PROB_AGING since $T_{last\_prob}$ , and update $T_{last\_prob}$ every time TX_Prob changes.

Note: MAX\_RETRY is the IEEE 802.11 MAC parameter for Maximum Retry Limit.

TX\_Prob value. Automatically increasing TX\_Prob after a certain time does not hurt the system performance; if the station is still disconnected, then TX\_Prob is reduced again when the first probing frame fails. The current default sets TX\_PROB\_AGING to 30 seconds. In many systems, an inactive TCP is terminated after 2 minutes. With 30 seconds for TX\_PROB\_AGING, it takes TX\_Prob at most 2 minutes to reach 1.0.

## 5 Simulation Results

### 5.1 Simple Scenario

In this section, the results show how a system reacts in the presence of the Bad Apple problem, both with and without SBA. The system model is the same as the one depicted in Figure 1. Simulation parameters are the same as those used in Section 3.

Figure 6 presents simulation results to illustrate how the Bad Apple flow affects the Victim flows, the Background flows, and the overall system performance. Goodput with time and instantaneous queue length (QLEN) are used to show what happens when the Bad flow is disconnected.

Figure 6 shows the results from two tests. The first test reproduces the BA problem in an 802.11 environment (without SBA), and its results are shown in the left column. The right column shows the test with SBA in the same system configuration. As shown in the left column, when the Bad flow's destination is unreachable for 10 seconds start-

ing at time 30, the Bad flow's goodput drops to 0. At the same time, due to the BA problem, the Victim flow's goodput drops significantly as well. However, there is no effect on the Background flow. The overall system goodput drops. The queue starts to fill and frame losses occur.

With SBA in the right column, the Bad flow has minimal impact on the Victim flow and the Background flow. As a result, the system-wide channel utilization is fairly stable. There is only a small spike in the queue length at 30 seconds before the SBA algorithm reacts.

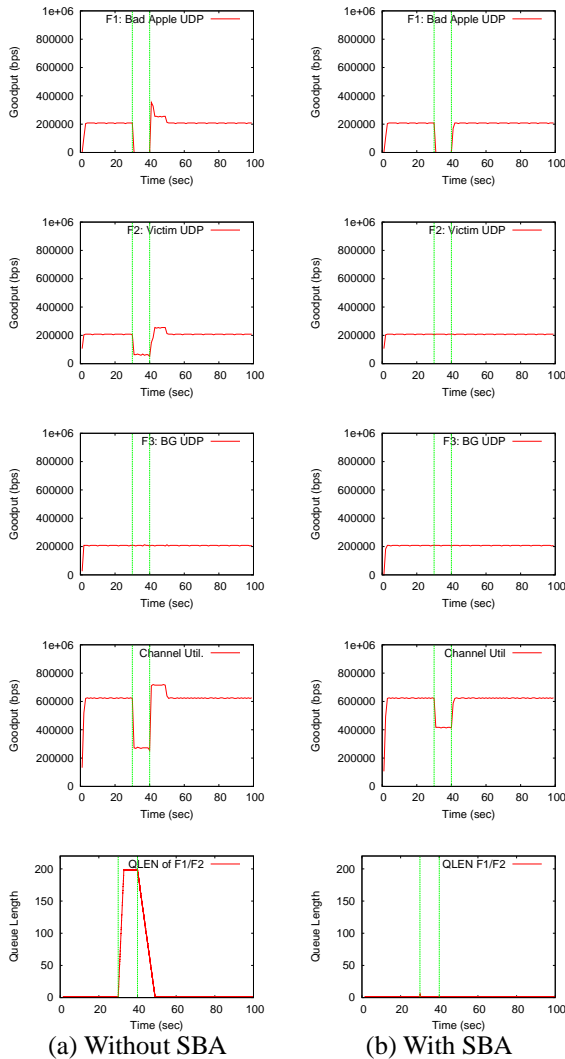
This simple test shows that SBA improves the robustness of IEEE 802.11. Moreover, the bandwidth freed by disabling the Bad flow can benefit other flows in the system. In the next section, we evaluate SBA in more complex scenarios.

### 5.2 Complex Scenarios

#### 5.2.1 Simulation Setup

This section evaluates the SBA approach in different system situations. We set up the system similar to the network shown in Figure 1 but with two UDP background flows. Again, F1 is the Bad flow, F2 is the Victim flow, and both are UDP flows. The MAX\_RETRY value is 7 for short frames and 4 for long frames. The channel capacity is set to 11 Mbps. The aggregate load is about 75%<sup>1</sup> of the effective bandwidth.

<sup>1</sup>Our previous experience [6] shows that when the WLAN is overloaded, the performance degrades for all clients. Therefore, we assume



**Figure 6. Illustration of Bad Apple Problem and Effect of SBA (F1,F2,F3: UDP 200 kbps)**

In all the tests, the connectivity to station C1 has alternating on-off periods. A short “off” period in wireless networks could be caused by walking behind a wall or having some temporary interference. A longer “off” period represents a user travelling out of the signal range. The “on” and “off” periods both follow exponential distributions. The “on” periods have a mean value of 10 seconds in all tests. The “off” periods have a mean of 10 seconds in Section 5.2.2, and different values in Section 5.2.3. All

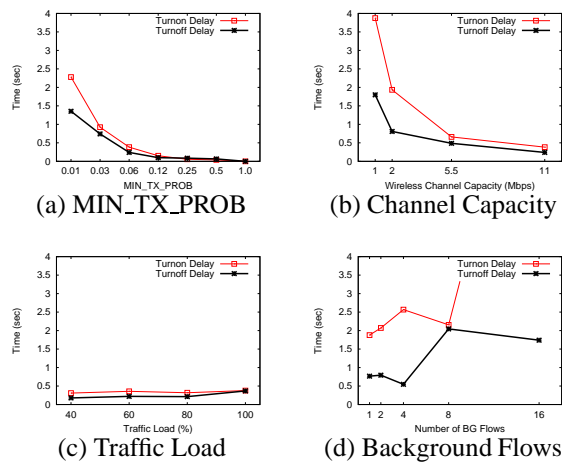
an effective admission control mechanism so that the system is never saturated. However, we have still studied how SBA reacts to system load and number of stations. Results in Figure 7(c) and (d) show that SBA is not overly sensitive to these settings. Based on these results, we use 75% load in all the experiments in this section.

tests use a long run of 1000 on-off cycles.

We first focus on deactivation delay and reactivation delay, and then show the overall system performance using average goodput, frame drop ratios, and average queueing delay as the performance metrics.

### 5.2.2 Deactivation/Reactivation Delay

The key factor that affects the system performance is how quickly the Bad flow can be disabled. In SBA, `TX_Prob` must be reduced to `MIN_TX_PROB`. We call this time the “deactivation delay”. The shorter the deactivation delay is, the quicker the bad flow is deactivated. We define the “reactivation delay” as the time from when the channel resumes to the time when `TX_Prob` reaches 1.0. The shorter the reactivation delay is, the sooner a flow can be resumed.



**Figure 7. Effects of Different Parameters on Deactivation/Reactivation Delays**

Figure 7(a) shows that the deactivation delay is inversely related to `MIN_TX_PROB`. That is, the lower `MIN_TX_PROB` is, the longer the deactivation delay is. The reactivation delay behaves similarly, but is slightly higher. When `MIN_TX_PROB` is 1.0, the deactivation delay is negligible, but as mentioned before, the Bad flow still wastes the channel. If `MIN_TX_PROB` is as low as 0.01, the Bad flow is almost disabled, but the deactivation delay could be as high as 2 seconds.

When `MIN_TX_PROB` is 0.06, the deactivation and reactivation delays are both less than 0.5 seconds. Therefore, we choose 0.06 as the default value for `MIN_TX_PROB`.

The commonly used 802.11b networks support multiple channel rates, namely 11 Mbps, 5.5 Mbps, 2 Mbps, and 1 Mbps. We test each channel capacity with an aggregate traffic load representing 80% of its effective bandwidth.

Figure 7(b) shows the deactivation and reactivation delays for different channel capacity settings. The lower the channel capacity is, the higher the delay is. The reason is that with low transmission rates, it takes longer to send a frame, and therefore longer for SBA to detect the disconnection. This result implies that SBA may not perform well in low bandwidth environments. Extreme care is required in multi-rate environments.

Figure 7(c) shows the simulation results for different traffic load levels (40%, 60%, 80%, and 100%), all with UDP flows. In all scenarios, the deactivation and reactivation delays are less than 0.5 seconds. These delays are not sensitive to the traffic load level. As long as the system is not overloaded, SBA can react quickly.

Since an 802.11 network is a shared broadcast environment, having additional active stations means less bandwidth share per station. Moreover, increased competition for channel access introduces higher variability to each flow’s transmission. We increase the number of background TCP flows to observe the effects of contention. Figure 7(d) shows that when the system is overloaded by TCP flows, the deactivation delay increases to about 2 seconds.


### 5.2.3 Performance with Longer-Term Dynamics

This section shows how SBA affects system performance in more complex scenarios. Four UDP flows (Bad, Victim, Background1, and Background2) are used, each sending 500 kbps. Good system performance means maintaining goodput for the Victim flows and limiting the queueing delay and frame drops (i.e., protecting other flows from the Bad flow).


We consider longer term simulations, with dynamic on-off behaviour for the Bad Apple flow. In each run, the “off” durations are drawn from an exponential distribution with a specified mean (0.1, 1, 5, 10, and 100 seconds). The brief 0.1-second “off” periods model a network with poor channel conditions at fine-grain time scales [5]. The “on” cycles are exponentially distributed with a mean of 10 seconds. Average goodput, average queueing delay, and frame drop ratios are used as the metrics to compare the performance with and without SBA.

Figure 8 (a) shows the average goodput with and without SBA. SBA does not affect the goodput of the Background flows. Also, in the cases of very short “off” periods (e.g., 0.1 second and 1 second), the channel utilization (aggregate goodput of all flows) with SBA is no worse than without.

The advantages of SBA are obvious when the “off” periods are longer. Without SBA, the goodput for the Victim flow (shaded) decreases due to the Bad Apple problem. With SBA, the Victim flow maintains its goodput. SBA outperforms previously proposed algorithms [5, 10] at longer time scales, and is no worse at shorter time scales.



(a) Goodput vs. Average Off Period



(b) Queueing Delay vs. Average Off Period

**Figure 8. Goodput and Queueing Delay with Longer-Term Dynamics**

Queueing delay results are depicted in Figure 8 (b). With longer time disconnections, without SBA, both the BA and the Victim flow experience high queueing delay (e.g., about 1 second when the “off” period is 10 seconds, and about 4 seconds when the “off” period is as long as 100 seconds). With SBA, queueing delays are about 0.001 seconds.

SBA also prevents frame drops. As can be seen in Table 4, with SBA, no queue overflow has been observed in any cases, whereas without SBA, there are many drops, especially when the “off” period is long.

**Table 4. Frame Drop Ratio Simulation Results**

Avg. On/Off Time (s)	Without SBA	With SBA
10 / 0.1	0.00%	0.00%
10 / 1	0.45%	0.00%
10 / 5	5.61%	0.00%
10 / 10	9.74%	0.00%
10 / 100	23.03%	0.00%

Maintaining high goodput, low queueing delay, and low drop ratios are crucial in 802.11 networks. If these basic requirements are not satisfied, QoS is impossible to achieve. The performance evaluation of SBA shows that it is a good



solution for the BA problem, providing robust support for QoS in IEEE 802.11 networks.

## 6 Conclusions

This paper studies the HOL blocking problem in IEEE 802.11 networks, and its manifestation as the Bad Apple problem. The BA problem can occur if multiple stations share a single FIFO transmission queue at the data link layer, the MAC-layer protocol employs retransmissions, and high-rate continuous UDP traffic is sent to a station that is unreachable for several seconds or more.

Our simulation results show that a Bad Apple flow adversely affects the victim flows. Based on our understanding of the BA problem, a Station-Based Adaptation (SBA) algorithm is proposed as an extension of IEEE 802.11 DCF. Simulation results show that the SBA mechanism can effectively protect the Victim flow from the BA phenomenon, improving the robustness of IEEE 802.11 networks.

Future work will investigate solutions for the BA problem in multi-channel multi-rate WLANs, since the BA problem can be particularly harmful in multi-rate environments. A combined strategy for both short-term and long-term disconnection will also be studied.

## Acknowledgements

Financial support for this research was provided by iCORE (Informatics Circle of Research Excellence) in the Province of Alberta, as well as NSERC (Natural Sciences and Engineering Research Council) in Canada, and TRILabs (Telecommunications Research Laboratories) in Calgary.

## References

- [1] G. Bai, K. Oladosu and C. Williamson, "Performance Issues for Wireless Web Servers", *Proceedings of 2004 International Workshop on Mobile and Wireless Ad Hoc Networking (MWAN)*, Las Vegas, NV, pp. 59-65, June 2004.
- [2] G. Bai and C. Williamson, "Simulation Evaluation of Wireless Web Performance in an IEEE 802.11b Classroom Area Network", *Proceedings of 3rd IEEE International Workshop on Wireless Local Networks (WLN)*, Bonn/Königswinter, Germany, pp. 663-672, October 2003.
- [3] G. Bai and C. Williamson, "The Effects of Mobility on Wireless Media Streaming Performance", *Proceedings of IASTED International Conference on Wireless Networks and Emerging Technologies (WNET)*, Banff, AB, Canada, pp. 596-601, July 2004.
- [4] A. Banchs, A. Axcorra, C. Garcia, and R. Cuevas, "Applications and Challenges of the 802.11e EDCA Mechanism: An Experimental Study", *IEEE Network*, Vol. 19, No. 4, pp. 52-58, July/August 2005.
- [5] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. Tripathi, "Enhancing Throughput over Wireless LANs using Channel State Dependent Frame Scheduling", *Proceedings of IEEE INFOCOM Conference*, San Francisco, CA, pp. 1133-1140, March 1996.
- [6] X. Cao, G. Bai, and C. Williamson, "Media Streaming Performance in a Portable Wireless Classroom Network", *Proceedings of IASTED European Workshop on Internet Multimedia Systems and Applications (EuroIMSA)*, Grindelwald, Switzerland, pp. 246-252, February 2005.
- [7] T. Henderson, D. Kotz, and I. Abyzov, "The Changing Usage of a Mature Campus-wide Wireless Network", *Proceedings of ACM MOBICOM*, Philadelphia, PA, pp. 187-201, September 2004.
- [8] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance Anomaly of 802.11b", *Proceedings of IEEE INFOCOM Conference*, San Francisco, CA, pp. 836-843, April 2003.
- [9] IEEE 802.11, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications", 1999.
- [10] L. Jiang and S. Liew, "An Adaptive Round Robin Scheduler for HOL Blocking Problem in WLANs", *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, New Orleans, LA, March 2005.
- [11] T. Kuang, Q. Wu, and C. Williamson, "MRMC: A Multi-Rate Multi-Channel MAC Protocol for Multi-Radio Wireless LANs", *Proceedings of the 2005 Workshop on Wireless Networks and Communication Systems (WiNCS)*, Philadelphia, PA, pp. 263-272, July 2005.
- [12] Q. Xia and M. Hamdi, "Cross Layer Design for IEEE 802.11 WLANs: Joint Rate Control and Packet Scheduling", *Proceedings of the 5th IEEE International Workshop on Wireless Local Networks (WLN)*, Sydney, Australia, pp. xx-yy, November 2005.