Experimental Evaluation of Optimal CSMA

Bruno Nardelli*, Jinsung Lee†, Kangwook Lee†, Yung Yi†, Song Chong†, Edward W. Knightly*, and Mung Chiang‡
*ECE Department, Rice University, Houston, TX
†Department of Electrical Engineering, KAIST, South Korea
‡Department of Electrical Engineering, Princeton University, Princeton, NJ

***Abstract*—By 'optimal CSMA' we denote a promising approach to maximize throughput-based utility in wireless networks without message passing or synchronization among nodes. Despite the theoretical guarantees on the performance of these protocols, their evaluation in real networking scenarios has been preliminary. In this paper, we propose a methodical approach for the first comprehensive evaluation of optimal CSMA, via experimentation with a custom implementation. Example findings include; 1) hidden terminals with symmetric channels can drive the protocol to a state of extreme contention aggressiveness due to the low service received by flows. Since increasing aggressiveness does not mitigate collisions but actually aggravates them, optimal CSMA enters a positive-feedback loop eventually reaching a deadlock state of total flow starvation; 2) however, the use of RTS/CTS in such scenarios can reduce collisions to lower levels, restoring throughput and preventing an excessive contention aggressiveness by optimal CSMA flows; 3) in practical hidden terminal scenarios with physical layer capture optimal CSMA reduces the aggressiveness of dominant flows, but the contention window sizes used by such adaptation mechanism are not long enough to solve competing flows' starvation when carrier sensing fails; 4) topologies with a "flow-in-the-middle" yield starvation in traditional CSMA but fairness in optimal CSMA, because its contention aggressiveness adaptation creates frequent transmission opportunities for the central (otherwise starved) flow; 5) optimal CSMA excessively prioritizes links with low channel quality, due to queue-based control that does not otherwise incorporate channel conditions; 6) in its current design, optimal CSMA conflicts with window-based end-to-end congestion control, and leads to a efficiency-fairness tradeoff in TCP performance. This study deepens our understanding of optimal CSMA and the general adaptation philosophy behind its design, and the derived insights suggest enhancements to optimal CSMA theory.**

I. Introduction

Recently, there has been an increasing interest from the research community in the design of distributed CSMA algorithms to maximize network utility [1]–[5]. The idea behind the operation of such protocols is to adapt the transmitters' contention aggressiveness as a function of flow queue lengths. As a link queue grows, the transmitter becomes more aggressive in the contention for channel access. The larger attempt rate increases its probability to access the channel as compared to competing flows. A remarkable aspect of these protocols is that they do not require any centralized control or message passing. In the following, we refer to this class of protocols as distributed optimal CSMA, or 'oCSMA' for short.

This work was supported by NSF grants CNS-1011962, CNS-1012831 and CNS-1012921, AFOSR grant FA9550-09-1-0643, ONR grants N00014-09-1-0449 and N00014-07-1-0864, the Cisco Collaborative Research Initiative, the Texas Instruments Leadership University Program, and IT R&D program of MKE/KEIT [KI002137, Ultra Small Cell Based Autonomic Wireless Network].

Under a list of assumptions, papers [1]–[5] presented proofs of convergence and optimality for oCSMA. Despite its simplicity and optimality, implementation and experimentation with oCSMA in practical settings have been very limited. The few existing implementations have been used only as a validation in simplistic scenarios [6], [7].

The key challenge in evaluating oCSMA is the design of modular experimental scenarios that reveal what aspects of oCSMA work or do not work and why. In addition, such scenarios must be realistic enough to extrapolate results to real-world operating conditions. To face this challenge, we propose a methodical approach, by which aspects critical to performance are isolated into carefully designed scenarios using a decoupling technique. Our technique provides insights into how each performance factor affects oCSMA, and enables the application of these results to oCSMA in real networks.

First, since network topology determines which nodes carrier sense each other, different interconnectivity among contending flows yield dramatically different throughput distributions in CSMA networks. For example, hidden terminals (HT) effectively disable carrier sense (CS) [8]–[11], whereas asymmetric connectivity can yield one or more nodes able to carrier sense while others cannot. Such a scenario, termed Information Asymmetry (IA) [10]–[12], can yield near starvation to the node unable to carrier sense. As a final example, uncoordinated interactions of neighbors may prevent a node from accessing the medium for long time periods, as in the Flow-in-the-Middle (FIM) topology [11], [13].

Second, the simultaneous occurrence of low-quality and high-quality channels poses problems of fairness and efficiency in wireless networks [14]. Furthermore, in topologies with hidden terminals, differences in the signal strength of concurrent transmissions make physical layer capture dominate over collisions [15], leading to unfair competition among flows comparable to that in topological asymmetries [10].

Finally, the execution of TCP congestion control over CSMA networks leads to throughput degradation [16], [17]. Moreover, severe starvation exists in scenarios with multiple competing TCP flows [18], [19]. Previous work has shown that the origins of such starvation are strongly related to properties of the MAC layer itself [20].

In the current oCSMA model, the above issues are mostly assumed away: perfect sensing, symmetric interference, homogeneous links, and infinite backlog not controlled by any window based upper layer. It remains to be seen if the predictive power of oCSMA theory remains despite these assumptions. More broadly, the new principle of MAC design in oCSMA: increase contention aggressiveness when being under-served, deserves a careful evaluation in realistic settings. By conducting the first comprehensive evaluation of oCSMA, we bring new discoveries that can drive the design of enhanced future oCSMA. Example findings are summarized below;

- With hidden terminals and symmetric channels, the inevitable collisions incurred by any CSMA variant cause oCSMA to become more aggressive due to lack of service. This high aggressiveness further increases collisions yielding a self-sustaining loop. Eventually, the protocol enters a state where no successful transmissions occur;
- The adaptation of contention aggressiveness by oCSMA is unable to solve the large throughput disparities due to information asymmetry. This is because the contention window sizes used in contention aggressiveness adaptation are too short to provide a significant improvement when transmissions are not coordinated by CS;
- In collision-prone scenarios such as topologies with hidden terminals and information asymmetry, the use of RTS/CTS allows oCSMA to attain better performance, by mitigating collisions and avoiding an excessive growth of flow contention aggressiveness.
- oCSMA solves problems of starvation in topologies with a flow-in-the-middle, which affect all other CSMA protocols with asynchronous uncoordinated transmitters. This is the first asynchronous CSMA solution to the FIM starvation problem validated through experimentation with real wireless hardware.
- oCSMA aims to assign the same throughput to all links regardless of the differences in their channel qualities. This is because oCSMA controls channel access priorities based only on queue lengths, without incorporating channel conditions.
- in its current design, oCSMA conflicts with upper-layer window-based congestion control, and leads to a efficiency-fairness tradeoff in TCP performance.

The remaining of this paper is organized as follows. Section II reviews background concepts on oCSMA. Section III describes our experimental methodology. Sections IV, V and VI respectively discuss the effects of topological factors, channel asymmetry and TCP congestion control on oCSMA networks. Finally, Section VII concludes.

II. Background: Optimal CSMA

As a distributed protocol, CSMA is easy to implement and has a low overhead, and is among the most widely implemented MAC protocols. In CSMA, a backlogged node waits for a random period of silent time before transmitting, termed *back-off time*. If no transmissions are sensed by the node (either by decoding headers of overheard packets or by measuring received energy) for the entire back-off period, the transmission starts. Otherwise, the node defers as soon as it senses an ongoing transmission, and resumes back-off after it senses the channel idle again. Back-off time distributions and transmission durations are two key factors that determine CSMA's dynamics and thus its efficiency. In practice, back-off times are implemented by discrete counters initialized to a random value of uniform distribution within a given *Contention Window* (CW).

TABLE I
Description of Optimal CSMA. $b[t]$: step size, $W(\cdot)$: weight function, and V, $q^{\min}$, $q^{\max}$ are positive parameters with $q^{\min} < q^{\max}$, and $[.]_d^c \equiv \max(d, \min(c, .))$.

1	**During each frame t** Run CSMA($\lambda_l[t], \mu_l[t]$), and record the amount of served packets $S_l[t]$ during this frame.
2	**At the end of each frame t** Update virtual queue q_l as $q_l[t+1] = \left[q_l[t] + \frac{b[t]}{W'(q_l[t])} \left(U'^{-1}\left(\frac{W(q_l[t])}{V}\right) - S_l[t] \right) \right]_{q^{\min}}^{q^{\max}}.$
3	Set $\lambda_l[t+1]$ and $\mu_l[t+1]$ such that their product is equal to $\exp(W(q_l[t+1]))$.

CSMA with *fixed* parameters may lead to poor performance due to either excessive collisions or too conservative channel access. Extensive attention has been paid to adaptive CSMA protocols (see e.g., [21] and the references therein) for high performance, where the key idea is to appropriately adapt back-off counters and/or transmission durations considering network conditions. Most protocols are developed from engineering heuristics and thus cannot guarantee a full coverage of the CSMA capacity region (802.11 DCF is a good example).

Recently, adaptive CSMA algorithms in multi-hop wireless networks have been revisited and generalized to target optimality in terms of throughput and fairness, using a utility maximization framework [1]–[5], [22]–[24]. More formally, oCSMA provably leads to the long-term link-level throughput which arbitrarily tightly solves the following optimization problem:

$$\max \quad \Sigma_{l\in\mathcal{L}} U(\gamma_l), \quad \text{such that} \quad \gamma \in \Gamma, \tag{1}$$

where $\mathcal{L}$ is the set of links and Γ is the set of *all possible* achievable rate vectors. The importance of this result is that oCSMA achieves optimality in this sense without requiring any explicit message passing or centralized schedulers.

We now explain the operation of oCSMA using the algorithm description in Table I. Assume that time is divided into successive frames. In Table I, CSMA(λ, μ) refers to CSMA having the random back-off counter with mean $1/\lambda$ and random transmission duration with mean μ. oCSMA maintains a virtual queue q_l for each link l, which keeps track of the amount of received service S_l. However, the virtual queue length is also affected by the amount of injected data, regulated by a congestion control mechanism. The utility function U is used to inject data at a rate that is inversely proportional to the virtual queue length (and such mechanism is termed *utility-based congestion control*). Finally, the back-off counter and the transmission duration are adapted as a function of the virtual queue length, so that an under-served link accesses the channel more aggressively than a well-served one.

Such an adaptation of contention aggressiveness minimizes the channel idle time, yet guarantees fairness among flows (at equilibrium, via the shape of the utility function U). However, the effectiveness of oCSMA and its underlying MAC design approach is proved under critical assumptions, such as perfect CS, symmetric interference, homogeneous links and infinite backlog not controlled by any window based upper layer protocol.

III. Experimental methodology

A. Topology

Our evaluation of oCSMA is based on the custom design of modular experimental scenarios. Such a modular problem decomposition enables a precise identification of which aspects of oCSMA work or not, and why. Furthermore, in real networks a combination of the studied situations can manifest simultaneously. Therefore, combining the conclusions that we derive for different experimental setups allows understanding the operation of oCSMA in a wide range of real-world operating conditions.

Fig. 1 shows the elemental scenarios that we use to study the effect of topological factors on the performance of oCSMA. First, this includes a fully-connected (FC) topology that we use as a baseline for comparison (Fig. 1a).

Second, a hidden terminal topology (HT) with symmetric channels (Fig. 1b). In this case, the operation of CS may provide incorrect information on the channel state, increasing the probability of simultaneous transmissions that collide at the receiver.

Third, in a topology with information asymmetry (IA, in Fig. 1c), concurrent transmissions over the two links collide at node 1 but not at node 3. This asymmetric situation leads to an unfair competition where the success rate of flow A is nearly zero while the success rate of flow B is nearly 1.

Finally, in a topology with a flow-in-the-middle (FIM, in Fig. 1d), the transmitter of the central flow can carrier sense transmissions over the side links (and vice versa), but the transmitters on the sides cannot carrier sense each other. Therefore, the central transmitter defers its transmissions whenever at least one of the side links is active, whereas concurrent transmissions on the side links can occur. Assuming no synchronization of nodes, transmissions on the side flows may interleave, leaving no silent periods for the central flow to start a new transmission. As the situation persists over time, the central flow starves while the side flows receive high throughput.

We evaluate the performance of oCSMA in each of these scenarios using fully-backlogged flows to magnify the effect of topological factors, in Section IV. Later, in Section VI, we use the same scenarios to study the joint effect of higher-layer congestion control and topological factors, with special attention to the FIM scenario.

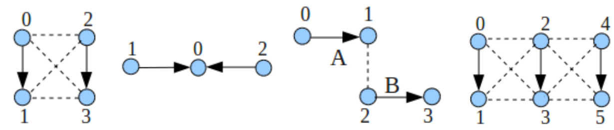

(a) FC (b) HT (c) IA (d) FIM

Fig. 1. Atomic topologies used to separate topological factors and study them in isolation. In the diagram, vertices represent network nodes, dotted lines represent the ability of nodes to carrier sense each other, and arrows represent traffic flows. Nodes are labeled by numbers and traffic flows are labeled by letters. Depicted topologies are; a. Fully Connected (FC); b. Hidden Terminals (HT); c. Information Asymmetry (IA); d. Flow-In-the-Middle (FIM).

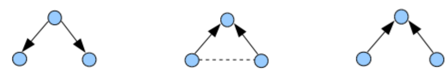

(a) AP-DL (b) AP-CS (c) AP-HT

Fig. 2. Atomic topologies used to separate factors related to channel asymmetry. In the diagram, vertices represent network nodes, dotted lines represent the ability of nodes to carrier sense each other, and arrows represent traffic flows. Depicted topologies are; a. an access point with two clients and downlink traffic (AP-DL); b. an access point with uplink traffic and clients in CS range of each other (AP-CS); c. an access point with uplink traffic and clients outside the CS range of each other (AP-HT).

B. Channels

Fig. 2 shows the atomic scenarios that we use to study the effects of channel asymmetry on the performance of oCSMA. First, when all flows share the same transmitter, there exists no physical contention among flows for channel access, but only a virtual contention implemented as a local decision at the transmitter node itself. This represents a best case for flow channel access prioritization, since the contention among flows is entirely controlled by the decision of a single node. To study this, we use the topology in Fig. 2a, with a transmitter and two receivers, which we call access point (AP) and clients, respectively. Given the similarities with a hotspot where users connect to download content, we call this scenario 'downlink AP case', for short AP-DL.

Second, when transmitters are at different nodes but in CS range of each other, collisions and physical layer capture are still rare, and most packet losses are due to channel errors. As oCSMA assigns higher priority to flows experiencing less throughput, it is likely that the presence of under-performing links will imply a reduction in the total network throughput. To study this case, we use the topology in Fig. 2b, with an AP and two clients in CS range of each other, which we refer to as AP-CS.

In a scenario with HTs, (see Fig. 2c), the larger the difference in signal strength among flows, the higher the probability of physical layer capture at the receiver. The extreme case manifests as an asymmetric interference relation where only packets of one flow are received. We refer to this topology as AP-HT.

C. oCSMA implementation

An implementation of oCSMA based on Common Code Architecture (CCA, [25]) was presented in [7]. In our experimental evaluation, we use this implementation, which comprises both a Glomosim-based simulator and a protocol

(a) Fixed testbed node (b) Mobile testbed node

Fig. 3. Experimental testbed hardware.

TABLE II
EXPERIMENTAL SETUP

Transmission rate	2 Mbps
Packet size	1000 bytes
Traffic pattern	Concurrent, fully-backlogged flows
Weight function $W(x)$	x
V	200
Utility function	$\log(x)$

implementation over standard 802.11 hardware. The oCSMA implementation over 802.11 hardware uses a modified *madwifi* driver to adapt the CW of transmitters at the device level as required by the oCSMA algorithm. For our evaluation of oCSMA, we deployed the CCA-based implementation in two platforms with different hardware for cross-validation and to minimize platform-specific results (see Fig. 3).

IV. Topological Factors

As highlighted by previous studies [8]–[13], topological factors have a tremendous impact determining the performance of CSMA protocols at the network scale. Ideally, the theoretically-proven optimal contention aggressiveness adaptation of oCSMA should help improve performance even in the presence of adverse conditions. Optimal CSMA has been designed from a new perspective than traditional CSMA. For example, continued packet losses have the effect of increasing the CW of transmitters when Binary Exponential Back-off (BEB) is used. For oCSMA, the effect is actually the opposite as packet losses imply a reduction on the received service, and consequent queue growth and increased contention aggressiveness. Thus, the importance of this section dedicated to evaluate oCSMA in such critical scenarios.

We use both simulations and experiments with our implementation on wireless hardware to study the effect of topological factors. Unless otherwise stated, the simulation results presented in this section correspond to 50 runs of 60 seconds each, and the experimental results correspond to 20 runs of 60 seconds each. All figures present average results with 95% confidence intervals. The oCSMA parameters in use are detailed in Table II.

A. Symmetric contention: FC and HTs

In fully connected topologies (FC), every transmitter can sense ongoing transmissions (albeit with a small propagation delay), and thus packet collisions rarely occur. Therefore, FC topologies represent a best-case scenario for the operation of CS. Because of this, we use a FC topology as a baseline when evaluating the performance of oCSMA in other scenarios.

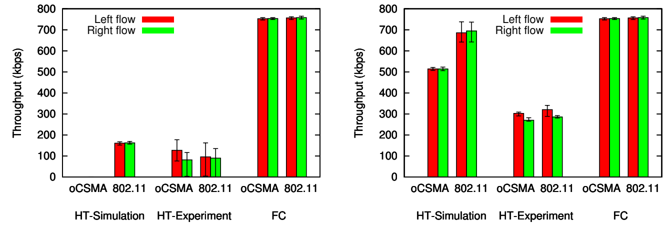

(a) RTS/CTS disabled (b) RTS/CTS enabled

Fig. 4. Per-flow throughput attained by 802.11 and optimal CSMA in a topology with hidden terminals. To provide for comparison, experimental and simulation results are presented together with those obtained in a fully connected topology.

It is well known that the presence of HTs can make CS fail, thus increasing the probability of collisions. Under such circumstances, protocols with BEB lower the attempt rate to time scales over the packet transmission time, which reduces the collision probability to moderate levels and improves throughput. RTS/CTS mechanisms, instead, attempt to move collisions from long data packets to shorter control packets which is again known to yield an improvement in performance.

In contrast, oCSMA interprets a growing queue length as being underserved compared to other flows and as a signal to increase aggressiveness in accessing the medium. But, in the HT case, the flows are both underserved (compared to the FC case) due to the ineffectiveness of CS. Consequently, an increase of transmission aggressiveness occurs at both senders. This, in turn increases even more the probability of collision worsening the problem. This interaction leads to a positive-feedback loop of increasing contention aggressiveness and increasing probability of collision. Eventually, the protocol reaches a deadlock state of maximum aggressiveness where no transmission can succeed.

Referring back to theoretical models, typical oCSMA designs assume no collisions. This is the assumption allowing it to infer that, in case of large queues, the node is being underserved relative to other flows. In practical scenarios with HTs, frequent collisions may affect both flows, and increasing the aggressiveness only worsens the problem.

However, the use of RTS/CTS mechanisms in this scenario improves the throughput of oCSMA flows. Not only the implied reduction in collision probability has the effect of increasing the success rate, but also to avoid entering the aforementioned loop of increasing contention aggressiveness.

Our simulation results in Fig. 4 support these arguments. The throughput of 802.11 flows is reduced to approximately a fourth in the HT scenario with respect to the FC case. Optimal CSMA flows, instead, completely starve with nearly zero throughput. In fact, the use of a simulator yields a strict condition of symmetric contention, where simultaneous transmissions always collide at the receiver. By the use of RTS/CTS, the situation can be significantly improved for both 802.11 and oCSMA, restoring flow throughput.

Experimenting in real wireless environments, instead, variable channels generate differences in the signal strength of simultaneous transmissions at the receiver. Even with very slight differences in the signal strength, physical layer capture intermittently allows packets to get received [15]. This results in a moderate-to-low throughput depending on the channel fluctuations, although the increase in contention aggressiveness is still observed in our experimental traces.

Findings: oCSMA's philosophy of increasing contention aggressiveness with flow queue length is inefficient in scenarios with high collision probability, such as topologies with hidden terminals. In such conditions, a symmetric increase of contention aggressiveness at both flows worsens the problem leading to a self-sustaining loop of performance degradation. RTS/CTS helps restoring oCSMA throughput, however to lower levels than for CSMA with BEB.

B. *Information Asymmetry (IA)*

A key scenario arising in practice is when symmetry is broken and one flow interferes with another but not vice versa. An example topology in which this occurs is depicted in Fig. 1c, and is referred to as the *Information Asymmetry* topology in [11].

In the IA topology, with CSMA, packets collide at receiver 1 but not at receiver 3. Consequently, when the network load is high, the success rate of the disadvantaged flow A can collapse to nearly zero while the success rate of the advantaged flow B still remains close to one. As opposed to the HT case, BEB does not help here, since only flow A experiences losses, and therefore only the disadvantaged transmitter 0 increases its CW size. Even with the use of RTS/CTS, IA can lead to large throughput disparities due to the same reasons [11].

In contrast, with oCSMA the high success rate in flow B keeps the queue length at the transmitter 2 low. This, in turn, implies that oCSMA should keep the CW of transmitter 2 at relatively large values. Ideally, such longer back-off times in the advantaged flow should provide additional silent time for the disadvantaged flow to successfully complete transmissions, leading to a more even distribution of throughput among flows.

However, in practice the use of oCSMA in IA topologies may lead to large throughput disparities comparable to the ones for 802.11 (see Fig. 5a). This is because in the IA scenario transmitter 0 is unable to determine when hindering transmissions are taking place. Therefore, its transmissions are only successful if, by chance, lie entirely within a silent period of the other flow. Thus, it cannot take complete advantage of generated transmission opportunities, and even the relatively long back-off times attained by the advantaged transmitter 2 with oCSMA are insufficient to significantly raise the success rate of the disadvantaged flow.

Enabling RTS/CTS, instead, the medium is reserved using a short control packet before each data transmission. Given the reduced length of an RTS packet, the probability of successful transmission considerably increases at the disadvantaged flow. This, together with the operation of contention aggressiveness adaptation of oCSMA, significantly raises the throughput of the disadvantaged (otherwise starving) flow.

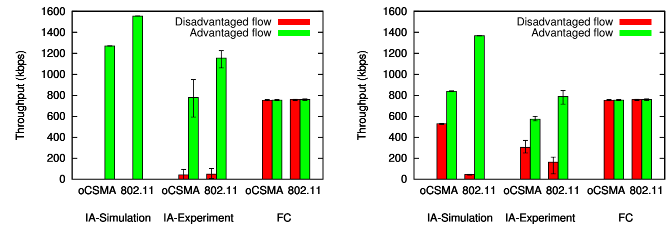

(a) RTS/CTS disabled (b) RTS/CTS enabled

Fig. 5. Per-flow throughput attained by 802.11 and optimal CSMA in a topology with information asymmetry. To provide for comparison, experimental and simulation results are presented together with those obtained in a fully connected topology.

We find supporting evidence in our experimental traces. Without RTS/CTS, the CW at transmitter 2 was always about 5 ms long (255 slots, 8 times larger than the typical CW size at the transmitter 0). Since back-off times are uniformly chosen within the CW, the average back-off at node 2 is close to 2.5 ms. With a packet length of about 4 ms (1000 bytes at 2 Mbps), the probability of transmitter 0 to complete a successful transmission is very low. However, since RTS control packets are much shorter, the RTS/CTS mechanism proves successful in solving this problem for oCSMA.

Findings: oCSMA by itself does not solve large throughput disparities due to information asymmetry. Even if oCSMA reduces the contention aggressiveness of the advantaged flow, the CW values used by contention aggressiveness adaptation are insufficiently long to provide a significant improvement. However, the joint operation of oCSMA with RTS/CTS overcomes this problem, and prevents flow starvation in topologies with information asymmetry.

C. *Flow-In-the-Middle (FIM)*

Problems in the previous scenarios are related to the imperfect operation of CS, which provides incomplete information to all transmitters (in the case of HTs) or to some of them (in the case of IA). Practical problematic scenarios can also arise from the sole interaction among transmitters, even with complete channel state information. This is the case of the topology depicted in Fig. 1d, referred to as Flow-In-the-Middle (FIM) in [11].

For a CSMA node to transmit, all of its neighbors must be silent. In the FIM topology, the side flows can transmit simultaneously, whereas the central flow can only transmit when both side flows are silent. Assuming no synchronization among nodes, silent times at the side flows may not coincide. As a result, the throughput of the central flow is lowered due to the lack of transmission opportunities, leading to complete starvation in the most extreme cases [11], [13].

This is a fundamental problem of any CSMA protocol with uncoordinated asynchronous transmitters. In fact, it affects all known protocols in this category, ranging from p-persistent CSMA to 802.11 variants, with and without the use of RTS/CTS mechanisms.

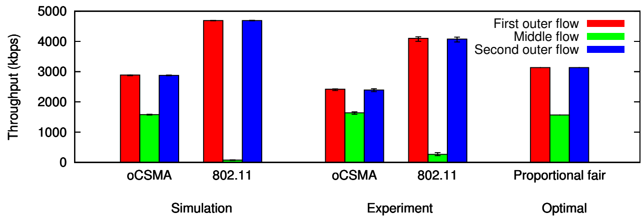

Fig. 6. Per-flow throughput attained by 802.11 and optimal CSMA in a topology with a flow-in-the-middle. To provide for comparison, the obtained results are presented together with the optimal throughput allocation for proportional fairness assuming an effective channel capacity of 4.7 Mbps. Simulation and experimental results were obtained using a 6 Mbps transmission rate.

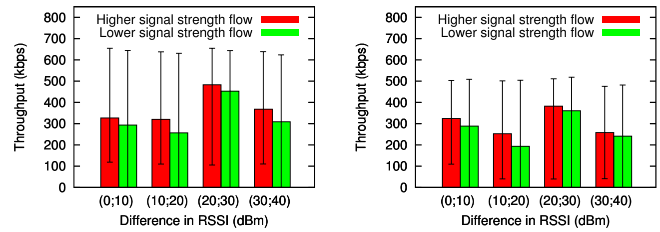

(a) 802.11 (b) Optimal CSMA

Fig. 7. Downlink flow throughput attained by two clients associated with an AP, in relation to the difference in RSSI measured at each of them.

In contrast, with oCSMA, the side transmitters that are in principle advantaged, maintain shorter flow queues than the central transmitter. As a result, the contention aggressiveness adaptation mechanism assigns the side flows a larger CW to lower their priority in channel access. This considerably increases the chances of simultaneous silent times at the side flows, generating more transmission opportunities for the central transmitter. In addition, the central transmitter, whose flow queue is larger, maintains a shorter CW, and quickly takes advantage of any generated transmission opportunities. This yields extremely good results for oCSMA, which successfully solves the FIM problem with a throughput distribution close to the optimal (see Fig. 6). The importance of these results lies in the generality of the FIM problem that, as explained before, affects all traditional CSMA protocols with uncoordinated asynchronous transmitters.

Findings: Optimal CSMA solves the FIM problem. The key to this success is the basic operation of contention aggressiveness adaptation, which generates more transmission opportunities for the central flow by enlarging the back-off times at the side transmitters, and reducing the back-off time for the central one. This is the first experimental validation showing that oCSMA solves the starvation problem in FIM topologies.

V. Asymmetric channels, capture and fairness

BEB doubles the CW size upon each failed transmission. Under high contention, this policy has the effect of reducing the probability of packet collision. However, in the presence of lossy channels, it can penalize traffic flows with higher loss rates, delaying their access to the channel.

In contrast, with oCSMA, frequent packet losses have the effect of increasing the flow queue length. In turn, this is interpreted as a signal to increase the channel contention aggressiveness of the node. Thus, on average, oCSMA assigns higher access priorities to lower-quality links. While this may increase the service of disadvantaged flows, it reduces system efficiency compared to schedules that opportunistically take advantage of channel fluctuation. Ideally, oCSMA's adaptation mechanism should help mitigating starvation problems related to physical layer capture. However, as for the topological asymmetries in Section IV-B, the silent back-off times might be too short to make a difference in such a competition with incomplete channel state information.

We now evaluate oCSMA in scenarios with channel asymmetry and physical layer capture. In order to gain a precise understanding on how the basic principle of contention aggressiveness adaptation copes with the above issues, we purposely focus on single rate scenarios, leaving modulation rate adaptation out of scope of this study.

A. *Fairness in lossy channels*

In a scenario where flows share the same transmitter, such as the one in Fig. 2a, CSMA with FIFO discipline serves packets in the order of arrival, irrespective of the flows they belong to. Also, most packet losses only imply a retransmission delay as opposed to dropping the head-of-line packet. Since all flows share the same queue, such additional delays affect all of them equally, and cause the same reduction in throughput. Then, if the flows sharing the queue have the same input rate, they all attain a similar throughput. Note that such an even throughput distribution is not proportional-fair, and makes the throughput of high-quality links disproportionately degrade in the presence of poor quality links.

In contrast, oCSMA senders use multiple queues, which allows control of channel access by different flows separately. Paradoxically, oCSMA suffers the same problem as CSMA in the presence of channel asymmetry with virtual contention. In oCSMA, channel access priorities are determined as a function of queue length. When multiple flows share the same sender, this mechanism is implemented as a deterministic decision equivalent to a 'longest-queue-first' policy. Furthermore, when the flow input rates are equal, such a policy leads to a round robin service of flows, where the transmitter sequentially serves a packet from the least-served-queue, retrying it until successful. Therefore, oCSMA flows sharing the same sender also attain a very similar throughput.

The above conclusions are derived from experimental results conducted in a network testbed with the topology in Fig. 2a. In total, we executed 90 experiment runs in 23 different client positions. We classify runs into 4 categories according to the difference in the mean RSSI in the signals over the two links (for the Atheros cards in our experimental testbed, RSSI is signal strength in dBm above the noise floor). The results obtained in terms of throughput for each of these categories are shown in Fig. 7. There, bars show average results with 95% confidence intervals. All other parameters are identical to the ones used in the previous section, detailed in Table II. Note the similarities between oCSMA and 802.11 in terms of throughput distribution, by which the two flows attain similar throughput irrespective of the differences in channel quality among them.

In scenarios with different transmitters in CS range of each other, such as the one in Fig. 2b, CSMA limits collisions and physical layer capture by the use of CS. With packet losses, BEB doubles the CW size at each retransmission. Then, transmitters over low-quality channels are penalized with longer mean back-off times than transmitters over high-quality channels. This amplifies the differences in throughput among flows due to channel errors.

In contrast, oCSMA adapts the contention aggressiveness of transmitters according to the received service. Since shorter back-off times are assigned to flows with longer queues, oCSMA prioritizes the access to the channel by underserved flows. In practice, such a least-served-first strategy leads to a sort of max-min fairness, which in the case of fully-backlogged flows targets an even distribution of throughput among flows.

The above conclusions are derived from experimental results conducted in a network testbed with the topology in Fig. 2b. In total, we executed 86 experiment runs in 22 different node positions. As we did for the previous case, we classify runs according to the difference in mean RSSI among links into 4 categories of 10 dBm each. The average results in terms of throughput are shown in Fig. 8 with 95% confidence intervals. 802.11 amplifies throughput disparities due to asymmetric channels whereas oCSMA attains an even distribution of flow throughput irrespective of the difference in channel quality among links.

Such an even throughput distribution does not correspond to the theoretical objective expressed by equation 1 in Section II. The reason for this contrast between theory and practical results is because the design of oCSMA assumes all links to have the same fixed capacity. Thus, oCSMA does not weight its decisions when controlling the access of multiple flows to asymmetric channels.

Findings: In scenarios with channel asymmetry, oCSMA targets an even distribution of flow throughput, as opposed to a proportional-fair distribution where each link attains a throughput proportional to its capacity. The reason is that oCSMA controls contention aggressiveness based only on queue length, without incorporating channel conditions.

B. Capture, asymmetries and starvation

A key scenario that arises in practice is when transmitters are out of CS range of each other, thereby increasing the probability of simultaneous transmissions. Such a scenario has been considered in Section IV-A with constraints of symmetric interference. But in CSMA networks it is often the case that the scenarios with HTs are dominated by physical layer capture instead of collisions [15]. Furthermore, the probability

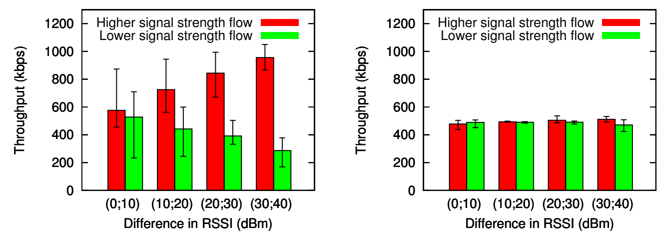

(a) 802.11 (b) Optimal CSMA

Fig. 8. Uplink flow throughput attained by two clients in CS range and associated with the same AP, in relation to the difference in RSSI measured from their signals at the receiver.

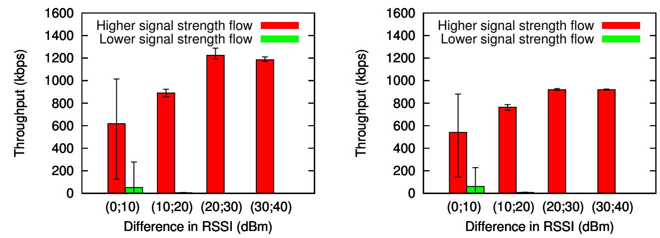

(a) 802.11 (b) Optimal CSMA

Fig. 9. Uplink flow throughput attained by two clients associated with the same AP but hidden from each other, in relation to the difference in RSSI measured from their signals at the receiver.

of physical layer capture over collision rapidly increases with the difference in signal strength among flows.

In extreme situations of physical layer capture, where one flow wins over the other with high probability, BEB has the effect of amplifying throughput disparities among flows. Even though BEB should in principle reduce the probability of simultaneous transmissions that overlap at the receiver, here it does not produce a significant benefit, but only adds delays to the disadvantaged flow in the channel access.

The effect of asymmetric interference induced by physical layer capture results in uneven growth of the flow queues for oCSMA. The advantaged flow, with high transmission success rate, maintains a much shorter queue than the disadvantaged flow, with a high loss rate and low throughput. In turn, the differences in queue length make oCSMA assign a much larger CW to the transmitter of the advantaged flow, while keeping the CW of the other transmitter short. Ideally, increasing the back-off times at the advantaged flow in this way should leave long silent periods for the disadvantaged flow to successfully complete transmissions.

However, this is not the case in practice, for the same reasons than in topological asymmetries discussed back in Section IV-B. The disadvantaged transmitter has incomplete information, and engages in a competition where its transmissions are only successful if, by chance, lie entirely within a silent period of the other flow. Thus, it cannot take complete advantage of generated transmission opportunities, and even the relatively large back-off times attained by the advantaged transmitter are insufficient to significantly raise its success rate.

The above conclusions are derived from experimental results conducted in a network testbed with the topology in Fig. 2c. In total, we executed 32 experiment runs. Again, we classify runs into 4 categories of 10 dBm each, according to the difference in the mean RSSI in the signals over the two links. The average results in terms of throughput are shown in Fig. 9 with 95% confidence intervals.

Findings: In scenarios with hidden terminals and asymmetric channels, the effects of channel capture can lead to the starvation of oCSMA flows, comparable to other CSMA protocols. This is because the CW values used in contention aggressiveness adaptation are too short to provide a significant improvement when transmissions are not coordinated by CS.

VI. Coupling oCSMA and End-to-end Congestion Control

In Section II and related work [26], oCSMA is described with the utility-based congestion control (i.e., the $U'^{-1}(\cdot)$ part in Table I and UBC for short) to highlight that in theory their joint operation provably guarantees optimality. However, TCP is a dominant congestion control protocol in practice, thus it is important to study the network performance when oCSMA is coupled with TCP.

Table III summarizes the experimental setup. We have tested four TCP variants, but present the results only for TCP Reno due to space limitation, its popularity in practice, and similarity of results compared to the other three variants. Unless otherwise specified, we use 16 kB $\text{CWND}_{\max}$ and 1024 byte MSS (Maximum Segment Size) for all experiments. Each data point is an average of 20 experiments with 95% confidence interval, each lasting for 60 seconds.

TABLE III
Setup for TCP Experiments

PHY	802.11a, 5.805 GHz band, 6 Mbps rate
TCP versions	Reno, Tahoe, NewReno, SACK
TCP MSS	512, 1024 bytes
TCP $\text{CWND}_{\max}$	16, 64, 256 kB
Step size, Weight function	$b[t] = 0.01$, $W(x) = x$

TCP is a window-based mechanism, where transmissions are clocked by acks. Its rate is controlled by AIMD (Additive Increase Multiplicative Decrease), the bounded number of in-pipe packets (i.e., $\text{CWND}_{\max}$), and large dependence of its throughput on RTT (Round Trip Time). In contrast, oCSMA dynamically adapts contention aggressiveness using the queue sizes and leverages congestion under heavy load to assign links' throughput close to the optimal point. This mismatch in their rationales causes oCSMA to miss its full potential when interacting with TCP. oCSMA's preference for large queue has also adverse impact on TCP's throughput, because TCP's throughput largely depends on RTT, i.e., CWND/RTT, and increasing queue size results in increasing RTT. This conflicting response to congestion causes the inter-play between the two to be non-trivial.

To demonstrate, consider the FC scenario in Fig. 1a. Fig. 10 shows how oCSMA+TCP behaves for bounded $\text{CWND}_{\max}$

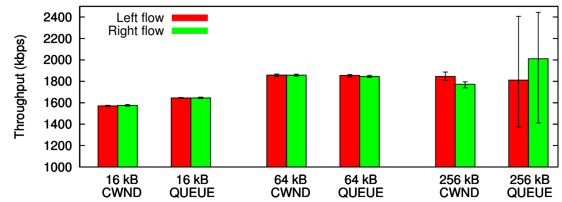

Fig. 10. Simulation. Effect of max physical queue size and $\text{CWND}_{\max}$ on TCP flows: Either bounded backlog by MAC or $\text{CWND}_{\max}$ by TCP can reduce TCP throughput since wireless channel is under-utilized due to increased back-off time with insufficient queued packets.

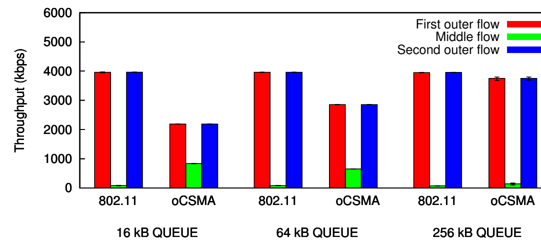

Fig. 11. Simulation. Per-flow throughput of 802.11 and oCSMA with different queue sizes in FIM topology.

and physical queue size. We observe that the channel is under-utilized for small values of $\text{CWND}_{\max}$ and physical queue size. This under-utilization is expected to be exacerbated in multi-hop flows because the maximum in-pipe packets may be spread over multiple queues along the path.

In Section IV-C, we verified that UBC+oCSMA achieves fair allocation of channel resource in FIM. In this section, we show that TCP+oCSMA regenerates unfairness in FIM which has complex coupling with RTT and bounded contention aggressiveness. There is a tradeoff between fairness and throughput in TCP+oCSMA. Fig. 11 shows that with 16 kB buffer size, fairness is almost guaranteed, which, however, comes at the cost of under-utilization of channel. With 256 kB buffer size, the total throughput is increased, but the central flow suffers from severe starvation.

With a small buffer size, the channel is under-utilized due to bounded contention aggressiveness, thereby low competition for the channel leads to fairness. However, with a large buffer size oCSMA enters the regime that aggressiveness is sufficiently exploited, thus the total throughput increases. The problem of TCP+oCSMA is that oCSMA's efficiency may be partially manifested only by the flows with less contenders (e.g., the outer flows in FIM). In UBC, the fairness is guaranteed because less service of the flow in the middle leads to larger backlogs, since the rate control keeps injecting data to the queue. However, in TCP, less service in that flow hinders the queue from increasing due to ack-clocking, i.e., data is inserted to the queue when CWND increases on reception of acks. This significantly delays the flush-out of the packets in the queue, and thus RTT also increases. Small CWND and large RTT starves the throughput of the flow in the middle. This is in stark contrast to UBC+oCSMA (see Fig. 12). TCP flow's starvation in FIM is also reported in [27] based on NS-2

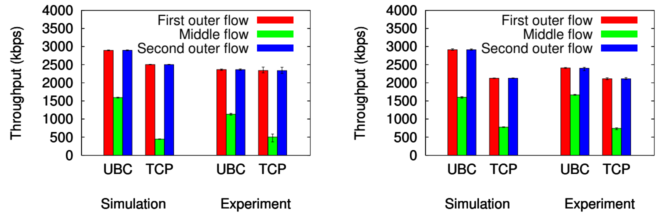

(a) 512 bytes packet and MSS. (b) 1024 bytes packet and MSS.

Fig. 12. Per-flow throughput of TCP and UBC in FIM topology: TCP $\mathrm{CWND}_{\max}$ is set to 16 kB, so that maximum backlogs can be attained by 32 (left) and 16 (right) packets for 512 (left) and 1024 (right) MSS, respectively.

simulations under a different context.

We have also performed experiments for HT and IA scenarios. The results are similar to those in UBC+oCSMA. An interesting observation in HT scenario is that one flow dominates over another flow even for *symmetric* links, in contrast to starvation of both flows in UBC+oCSMA. This is due to the fact that protocol-specific behavior and lack of synchronization inherent in TCP protocol eventually make one flow receive more service than the other, after which that flow monopolizes the channel. In experiments, due to capture effect, channel asymmetry is inevitable, and thus the channel monopolization is consistently seen.

Findings: Interaction with higher-layer window-based congestion control can have a severe impact on the throughput distribution among oCSMA flows. In networks with a flow-in-the-middle and small buffer size, the channel is under-utilized due to bounded contention aggressiveness, but the resource allocation is reasonably fair. In contrast, with large buffer size, the middle flow can starve completely due to small CWND and large RTT, which differs significantly from UBC+oCSMA.

VII. Conclusion

Our results highlight the need for oCSMA model to capture physically accurate collision scenarios, asymmetric interference, heterogeneous channels and realistic traffic patterns. This study feeds back the oCSMA design community with a better understanding on protocol operation in realistic settings, completing one round along the loop of model-design-implementation-data-model. Furthermore, the modular evaluation approach proposed here reveals what aspects of oCSMA work or do not work and why, suggesting directions for future design. Example possible directions include: reducing contention aggressiveness while increasing channel holding time, providing visibility across TCP-MAC boundary, and incorporating channel conditions to guarantee proportional fair throughput distribution in the presence of asymmetric channels.

References

[1] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," in *Proc. Allerton Conference*, Monticello, IL, 2008.

[2] J. Ni and R. Srikant, "Distributed csma/ca algorithms for achieving maximum throughput in wireless networks," in *Proc. ITA*, San Diego, CA, 2009.

[3] P. Marbach and A. Eryilmaz, "A backlog-based csma-mechanism to achieve fairness and throughput-optimality in multihop wireless networks," in *Proc. Allerton Conference*, Monticello, IL, 2008.

[4] S. Rajagopalan and D. Shah, "Distributed algorithm and reversible network," in *Proc. of CISS*, 2008.

[5] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor, "Towards utility-optimal random access without message passing," *Wiley Journal of Wireless Communications and Mobile Computing*, vol. 10, no. 1, pp. 115–128, Jan. 2010.

[6] A. Sridharan, S. Moeller, and B. Krishnamachari, "Implementing backpressure-based rate control in wireless networks," in *Proc. ITA Workshop*, San Diego, CA, 2009.

[7] J. Lee, J. Lee, Y. Yi, S. Chong, A. Proutiere, and M. Chiang, "Implementing utility-optimal CSMA," in *Proc. Allerton Conference*, Monticello, IL, 2009.

[8] F. A. Tobagi and L. Kleinrock, "Packet switching in radio channels: Part ii - the hidden terminal problem in carrier sense multiple access and the busy tone solution," *IEEE Trans. on Communications*, vol. 23, no. 12, pp. 1417–1433, 1975.

[9] P. Karn, "MACA: A new channel access method for packet radio," in *9th Computer Networking Conference*, 1990.

[10] V. Bharghavan, S. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless lans," in *Proc. ACM Sigcomm*, London, UK, 1994.

[11] M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," in *Proc. IEEE Infocom*, Barcelona, Spain, 2006.

[12] V. Bharghvan, "Performance evaluation of algorithms for wireless medium access," in *Proc. IEEE IPDS*, Durham, NC, 1998.

[13] J. Shi, E. Aryafar, T. Salonidis, and E. Knightly, "Synchronized CSMA contention: Model, implementation and evaluation," in *Proc. IEEE Infocom*, Rio de Janeiro, Brazil, 2009.

[14] M. H. Franck, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE Infocom*, San Francisco, CA, 2003.

[15] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Berlin, Germany, 2004.

[16] M. Gerla, K. Tang, and R. Bagrodia, "TCP performance in wireless multi-hop networks," in *Proc. IEEE WMCSA*, New Orleans, LA, 1999.

[17] K. Sundaresan, V. Anantharaman, H. Hsieh, and R. Sivakumar, "ATP: A reliable transport protocol for ad-hoc networks," in *Proc. ACM Mobihoc*, Annapolis, MD, 2003.

[18] K. Tang and M. Gerla, "Fair sharing of MAC under TCP in wireless ad-hoc networks," in *Proc. IEEE MMT*, Venice, Italy, 1999.

[19] K. Xu, M. Gerla, L. Qi, and Y. Shu, "Enhancing TCP fairness in ad-hoc wireless networks using neighborhood RED," in *Proc. ACM Mobicom*, San Diego, CA, 2003.

[20] J. Shi, O. Gurewitz, V. Mancuso, J. Camp, and E. Knightly, "Measurement and modeling of the origins of starvation in congestion controlled mesh networks," in *Proc. IEEE Infocom*, Phoenix, AZ, 2008.

[21] F. Cali, M. Conti, and E. Gregori, "Dynamic tuning of the ieee 802.11 protocol to achieve a theoretical throughput limit," *IEEE/ACM Trans. on Networking*, vol. 8, no. 6, pp. 785–799, 2000.

[22] J. Shin, D. Shah, and S. Rajagopalan, "Network adiabatic theorem: An efficient randomized protocol for contention resolution," in *Proc. ACM Sigmetrics*, Seattle, WA, 2009.

[23] M. Durvy and P. Thiran, "Packing approach to compare slotted and non-slotted medium access control," in *Proceedings of IEEE Infocom*, 2006.

[24] L. Jiang and S. Liew, "Improving throughput and fairness by reducing exposed and hidden nodes in 802.11 networks," *IEEE Trans. on Mobile Computing*, vol. 7, no. 1, pp. 34–49, 2008.

[25] J. Lee, J. Lee, K. Lee, and S. Chong, "Commoncode: A code reuse platform for co-simulation and experimentation," in *Proc. ACM CoNEXT Student Workshops*, Philadelphia, PA, 2010.

[26] J. W. Lee, M. Chiang, and R. A. Calderbank, "Jointly optimal congestion and contention control in wireless ad hoc networks," *IEEE Communications Letters*, vol. 10, no. 3, pp. 216–218, 2006.

[27] W. Chen, Y. Wang, M. Chen, and S. Liew, "TCP over adaptive CSMA," submitted for publication.