

Making 802.11 DCF Near-Optimal: Design, Implementation, and Evaluation

Jinsung Lee, Hojin Lee,
Yung Yi and Song Chong
EE Dept., KAIST, Korea

Bruno Nardelli
and Edward W. Knightly
ECE Dept., Rice Univ., USA

Mung Chiang
EE Dept. Princeton Univ.
Princeton, NJ, USA

Abstract—This paper proposes a new protocol called *Optimal DCF (O-DCF)*. O-DCF modifies the rule of adapting CSMA parameters, such as backoff time and transmission length, based on a function of the demand-supply differential of link capacity captured by the local queue length. O-DCF is fully compatible with 802.11 hardware, so that it can be easily implemented only with a simple device driver update. O-DCF is inspired by the recent analytical studies proven to be optimal under assumptions, which often generates a big gap between theory and practice. O-DCF effectively bridges such a gap, which is implemented in off-the-shelf 802.11 chipset. Through extensive simulations and real experiments with a 16-node wireless network testbed, we evaluate the performance of O-DCF and show that it achieves near-optimality in terms of throughput and fairness, and outperforms other competitive ones, such as 802.11 DCF, optimal CSMA, and DiffQ for various scenarios.

I. INTRODUCTION

Extensive research documents inefficiency and unfairness of the standard 802.11 DCF and suggests ways to improve it. Some proposals take a clean-slate approach to redesign CSMA. For some of them, optimality in performance can sometimes be proved under idealized assumptions such as no collision or perfect synchronization. Other solutions are constrained to operate over legacy 802.11 hardware with only a device driver update, but their performance improvements are often marginal. Among prior methods to improve 802.11 DCF is the seemingly conflicting pair of random access philosophies: in face of collisions, should transmitters become more aggressive given that the supply of service rate may become lower than the demand (as in the recently developed theory of Optimal CSMA (oCSMA), e.g., [1]–[5])? Or should they become less aggressive given that collisions signal a contentious RF environment (as in a typical exponential backoff in 802.11 DCF)?

In this paper, we propose a new protocol, called *Optimal DCF (O-DCF)*, showing that the two approaches are in fact complementary, where the best combination depends on the logical contention topology, but can be learned without knowing the topology. O-DCF is inspired on the ideas of CSMA adaptation developed by different analytical optimization work [1]–[5], but it is designed to (i) be fully compatible with off-the-shelf 802.11 chipsets, and (ii) achieve high performance in practice by tackling practical issues encountered in actual deployment, which are largely ignored in theory, yet have high impact on actual performance.

In O-DCF, a product of access probability, which is determined by contention window (CW) size, and transmission

length is set to be proportional to the supply-demand differential for long-term throughput fairness. Towards the goal of achieving high performance in practice, a good combination of access probability and transmission length is taken, where such a good access probability is “searched” by Binary Exponential Backoff (BEB) in a fully distributed manner to adapt to the contention levels in the neighborhood, and then transmission length is suitably selected for long-term throughput fairness. Thus, BEB is exploited not just to conservatively respond to temporal collisions, as in standard 802.11, but also to be adapted to appropriate access aggressiveness for high long-term fairness and throughput by being coupled with the transmission length.

We first summarize three key design ideas of O-DCF:

- D1.** Link access aggressiveness is controlled by both CW size and transmission length, based on per-neighbor local queue length at MAC layer, where the queue length quantifies supply-demand differential. Links with bigger differential (i.e., more queue buildup) are prioritized in media access by decreasing CW size and/or increasing transmission length.
- D2.** The CW size and transmission length are adapted in a fully distributed manner, depending on network topology affecting contention patterns in the neighborhood. Each link first chooses the initial CW size as a decreasing function of the local queue length and increases the CW size against collisions (i.e., BEB). Then, transmission length is decided by a “smart” combination of the resulting CW size (at which transmission succeeds) and the local backlog.
- D3.** When wireless channels are heterogeneous across links, e.g., a 2Mbps link and another 6Mbps link, the link capacity information is reflected in controlling access aggressiveness by scaling the queue length proportionally to the link capacity. This adaptive control based on link capacity ensures better fairness and higher throughput, since links with better channel condition are scheduled more than those with poor channel condition.

The key design ideas mentioned above are implemented through the following protocol mechanisms:

- PI.** Each transmitter maintains two queues for each neighbor, referred to as Control Queue (CQ) and Media Access Queue (MAQ). CQ buffers the packets from the upper-

layer which are to be dequeued into MAQ. The length of MAQ refers to the local queue length in **D2**, determining access aggressiveness by adjusting CW size and transmission length. The dequeue rate from CQ to MAQ is controlled through the MAQ length to ensure various, tunable fairness criteria and high throughput.

- P2.** Once the initial CW size is chosen as a function of the length of MAQ, BEB “searches” for the CW size at which transmission becomes successful in a fully distributed manner. This successful CW size is used to choose the transmission length as described in **D2**.
- P3.** We adapt transmission length based on time rather than bytes to achieve time fairness even under heterogeneous channels. To that end, we exploit instantaneous link information from the rate-adaptation module in the 802.11 driver to determine the proper number of bytes to send, according to modulation and coding rate in use.

All of the above mechanisms can be implemented using unmodified 802.11 chips, as we have done in evaluating O-DCF over a 16-node wireless testbed. In particular, the mechanisms satisfy the following constraints of operation over 802.11; first, 802.11 hardware uses a single queue (called Interface Queue (IQ)) shared by all local flows, so that neighbor-specific packet control necessitates additional queues such as CQ and MAQ on top of the IQ. Second, CW values are allowed by only some powers of two and thus we can only choose the value from the set $\{2^n - 1, n = 1, \dots, 10\}$. Third, the transmission length is bounded by some value that depends on the 802.11 chipset. For example, it is determined by the minimum of maximum allowable time and maximum aggregation size, e.g., 10 ms and 64 KB in 802.11n [6].

To evaluate the performance of O-DCF, we have implemented O-DCF on a 16-node wireless testbed as well as a simulator for large-scale scenarios that are difficult to be configured in the real testbed. By comparing O-DCF with 802.11 DCF, two versions of oCSMA, and DiffQ [7], we observe that in presence of conditions that are known to be critical to other CSMA protocols, O-DCF achieves near-optimal throughput, fairly distributed among flows, with up to 87.1% fairness gain over 802.11 DCF.

II. RELATED WORK

Numerous papers have reported the performance problems of 802.11 DCF, and proposed many solutions to them. To name just a few, 802.11 DCF has severe performance degradation and throughput disparities among contending flows in the topologies such as Hidden-terminal (HT), information asymmetry (IA), flow-in-the-middle (FIM), and packet capture [8]–[11], and heterogeneous link capacities [12]. We classify the solution proposals into the efforts at MAC and/or PHY. Some papers proposed new access methods such as dynamic adjustment of CW under 802.11 DCF, e.g., [13], [14] and there exist the implementation researches along this line, e.g., [15], [16]. Other work presented efficient aggregation schemes and their real implementations for throughput improvement, e.g., [17], [18].

Note that most implementations mentioned above individually focused on some specific topologies such as fully-connected (FC) case, and do not explicitly consider problematic ones such as HT, IA, and FIM. With the aid of PHY-layer, there are totally novel approaches, e.g., [19], [20]. This kind of work exploits more information from PHY layer and/or applies new PHY technologies other than CSMA. In O-DCF, we extend MAC-level operation *by controlling the CW size as well as the transmission length based on the demand-supply differential*. The PHY-layer based approach is somewhat orthogonal to our approach, which even can be integrated with O-DCF for further performance improvement.

Recently, analytical studies proved that, under certain assumptions, queue-length based scheduling via CSMA can achieve maximum throughput without message passing e.g., [1]–[3], which is referred to as oCSMA in this paper. Furthermore, multiple theoretical papers presented solutions based on the similar mathematical framework, each of them focusing on different aspects of the protocol operation [4], [5]. Our work is in part motivated by oCSMA theory, but as reported in [21], [22] and our evaluation, there still exist many gaps between oCSMA theory and 802.11 practice. There exists work that bridges the gaps between theory on queue-based MAC (not necessarily CSMA) and practice by reflecting queue length over 802.11 [7], [23], [24]. In [7], the authors implemented a heuristic differential backlog algorithm (DiffQ) over 802.11e. EZ-Flow was proposed to solve instability due to large queue build-ups in 802.11 mesh networks [23]. Very recently, the authors in [24] implemented a backpressure scheduling based on TDMA MAC, which operates in a centralized way. In contrast with the aforementioned work, we present the first asynchronous CSMA based protocol, motivated by optimal CSMA theory and designed to work on top of the legacy 802.11 hardware, that can attain high performance even in several adverse scenarios.

III. HOW DOES O-DCF WORK?

A. Balancing Supply-Demand Differential

In 802.11 DCF, the packets from the upper-layer are enqueued to IQ at the 802.11 chip for media access. In O-DCF, we additionally maintain two per-neighbor queues (CQ and MAQ) over IQ, to balance the link’s supply and demand through fair media access, as shown in Fig. 1. Denote by $Q_l^C(t)$ and $Q_l^M(t)$ the lengths of CQ and MAQ for each link l at time t . We further maintain a variable $q_l^M(t)$, which is simply the scaled version of $Q_l^M(t)$, i.e., $q_l^M(t) = bQ_l^M(t)$, where b is some small value.¹ Finally, the value of $Q_l^M(t)$ is crucial in O-DCF in that both the dequeue rate from CQ to MAQ and the aggressiveness of media access tightly rely on $Q_l^M(t)$.

First, we control the dequeue rate from CQ to MAQ (when CQ is non-empty), such that it is inversely proportional to

¹The value b is referred to as a step size in theory, which is used to slow down the variation of queue lengths, as well as to have feasible CW sizes in the aggressiveness update, see Remark 5.1.

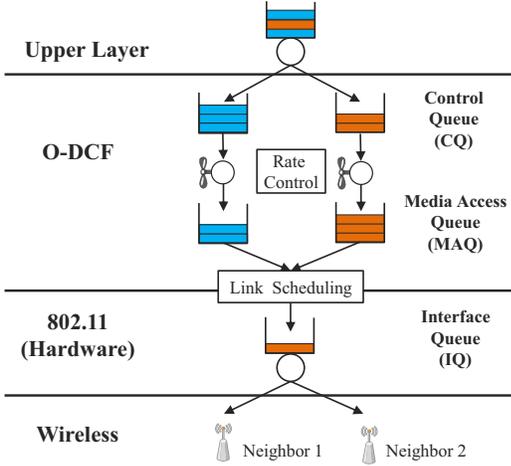


Fig. 1. Queue structure of O-DCF.

$q_i^M(t)$, by $V/q_i^M(t)$,² where V is some constant controlling the sensitivity of the dequeue rate to MAQ. Second, in terms of aggressiveness in media access, an initial CW size and transmission length, which determine the dequeue rate of MAQ, is set as a function of $q_i^M(t)$ (see Sections III-B and III-C for details). Then, whenever a new arrival from CQ or a service (i.e., packet transmission) from successful media access occurs, $Q^M(t)$ is updated by:

$$Q_i^M(t + \delta t) = [Q_i^M(t) + (\text{arrival from CQ} - \text{service from MAQ})]_{Q_{\min}}^{Q_{\max}}, \quad (1)$$

where δt is the elapsed time of the next arrival or service event after t . The service from MAQ occurs when the HOL (Head-Of-Line) packet of MAQ is moved into IQ. For multiple neighbors, the largest MAQ is served first; If the chosen transmission length exceeds a single packet size, multiple packets from the same MAQ are scheduled in succession. Q_{\max} is the physical buffer limit of MAQ, but Q_{\min} is set as a small positive value to prevent impractically high injection rate from CQ to MAQ, when $Q_i^M(t)$ (and thus $q_i^M(t)$) approaches zero.

B. Initial CW with BEB

When a link l is scheduled at time t , its initial CW size ($CW_l(t)$) and the number of bytes to transmit ($\mu(t)$) are set adaptively as a function of the length of MAQ. First, $\mu(t)$ -byte transmission over l is assigned with the following CW size:

$$CW_l(t) = \frac{2(\exp(q_i^M(t)) + C)}{\exp(q_i^M(t))} - 1, \quad (2)$$

where C is some constant whose suitable value will be discussed in Section IV-B. We want to use $CW_l(t)$ as the minimum CW (CW_{\min}) in 802.11 DCF. But since the 802.11 hardware only allows CW values as powers of two, We use

²This form of injection pattern is for achieving proportional fairness. However, as will be explained in Section IV-A, this can be changed towards other fairness criterion.

one of possible values closest to (2) as CW_{\min} . Then, the media access is attempted after the time (in mini-slots) randomly chosen from the interval $[0, CW_{\min}]$. Intuitively, we assign higher aggressiveness in media access for larger MAQ size, remarking that (2) is decreasing with $q_i^M(t)$. Whenever collision happens, this CW value increases exponentially by BEB. In Section IV-B, we will explain that BEB is not just a component often constrained by some legacy 802.11 hardware,³ but is an important component to choose an appropriate access aggressiveness and reduce collisions inside our design rationale in a distributed manner.

It is often convenient to interpret CW_l with its corresponding access probability p_l using the relation $p_l = 2/(CW_l + 1)$ [13], [14], where the initial CW_l selection in (2) is regarded as the following *sigmoid* function:

$$p_l(t) = \frac{\exp(q_i^M(t))}{\exp(q_i^M(t)) + C}. \quad (3)$$

We delay our discussion on why and how this sigmoidal type function helps and what choice of C is appropriate to Section IV-B.

C. Transmission Length Selection

Our method for the adaptation of the transmission length $\mu_l(t)$ upon channel access by link l is to set $\mu_l(t)$ as a function of the success access probability (equivalently, the success CW size) and the size of MAQ (i.e., $Q_i^M(t)$). Again, by the success CW size, we mean that the transmission becomes successful at that CW size, which is often larger than the initial CW size due to BEB. The rationale to search for the success CW size lies in the fact that it is the actual value used in media access for successful transmission. However, it usually requires quite high overhead for the device driver to read such a success CW size in real-time, which may result in delaying packet scheduling. Thus, we estimate it based on the equation [26] which describes the connection between the initial CW (CW_l), collision ratio (p_c), and the success access probability after BEB (denoted by \tilde{p}_l), given by:

$$\tilde{p}_l = \frac{2q(1 - p_c^{m+1})}{(CW_l + 1)(1 - (2p_c)^{m+1})(1 - p_c) + q(1 - p_c^{m+1})}, \quad (4)$$

where $q = 1 - 2p_c$ and m is the maximum retransmission limit. The collision ratio p_c can be computed for arbitrary topologies [9], only if nodes have the complete knowledge of topology via message passing. To avoid such message passing, we measure the packet collisions by counting the number of unacked packets and normalizing them. Using $\tilde{p}_l(t)$, in O-DCF, the transmission length (in mini-slots) is chosen by:

$$\mu_l(t) = \min\left(\frac{\exp(q_i^M(t))}{\tilde{p}_l(t)}, \bar{\mu}\right), \quad (5)$$

where $\bar{\mu}$ is the maximum transmission length (e.g., 64 KB and 10 ms as mentioned earlier, in order to ensure the minimum

³In literature, e.g., [25], it is known that BEB can be disabled by using TXQ descriptor in the device driver, which, however, does not work in our chipset, confirmed via our kernel level measurement.

short-term fairness and prevent channel monopolization by some node). Then, we convert the transmission length in the unit of mini-slots of the 802.11 chipset into that in bytes to compute the number of packets for aggregate transmission by μ_l (bytes) = μ_l (slots) \times c_l (Mb/s) \times t_{slot} ($\mu\text{s}/\text{slot}$). When only a part of μ_l bytes is transmitted due to packetization, we maintain a *deficit counter* to store the remaining bytes of the transmission length that will be used in the next transmission.

D. Channel Heterogeneity and Imperfect Sensing

In practice, wireless channels are heterogeneous across users as well as often time-varying. In such environments, most 802.11 hardware exploits multi-rate capability of the PHY layer to adapt their rate, e.g., SampleRate [27]. However, it is known that 802.11 DCF is incapable of utilizing this opportunistic feature, leading to the waste of resource called *performance anomaly* [12]. In other words, 802.11 DCF provides equal chances to the links (on average), in which case the low-rate links would occupy more time than the high-rate ones, so that the performance degrades. To provide fairness focusing on time shares instead of rate shares, namely *time-fairness* [28], we slightly modify our rules in selecting the initial access probability as well as the transmission length by replacing $\exp(q_l^M(t))$ with $\exp(c_l(t)q_l^M(t))$ in (2) and (5), where $c_l(t)$ is (relative) link capacity of link l at time t , as theoretically verified by [29]. For imperfect sensing cases such as HT and IA scenarios, we also propose to use a virtual carrier sensing via RTS/CTS signaling, as suggested in literature. In O-DCF, unlike in the standard 802.11a/b/g, RTS/CTS signaling is conducted only for the first packet within the transmission length.

IV. WHY DOES O-DCF WORK?

A. Transferring from Theory to Practice

O-DCF is in part motivated by the recent research on queue-based MAC scheduling in theory community, see e.g., a survey [30] and in particular, so-called oCSMA [1]–[5]. oCSMA is characterized as a variant of CSMA that has a specific rule of setting backoff time and transmission length. The papers in literature are slightly different in terms of the models and conditions, e.g., discrete/continuous, synchronous/asynchronous, or saturated/unsaturated traffic. However, the key idea is largely shared; the queue maintains the demand-supply differential, and the access aggressiveness is controlled by the queue length, which, in turn, depends on the demand (arrival) and the supply (transmission success), formally $p_l(t) \times \mu_l(t) = \exp(q_l(t))$. Together with this parameter control, the source rate control by $U'^{-1}(q(t)/V)$, where $U(\cdot)$ is an utility function, ensures that the long-term throughput is the solution of NUM (network utility maximization) problem. By suitably choosing the form of the utility function, we can achieve various fairness criteria. In this paper, we focus on $U(\cdot) = \log(\cdot)$ (thus, $U'^{-1}(q(t)/V) = V/q(t)$).

We highlight that O-DCF is not just a naive implementation of oCSMA, because (i) many assumptions in the oCSMA theory, e.g., no collisions in the continuous time framework, symmetric sensing, perfect channel holding and etc. do not



Fig. 2. Example topologies' conflicting graphs. Vertices are the links in interference graphs; dotted lines represent interference; (a) FC with 4 links; (b) FIM with 2 outer links.

hold in practice. (ii) Furthermore, O-DCF is constrained to be fully compatible with 802.11 chipsets. (iii) More importantly, in theory, any combination of p_l and μ_l works well as long as their product is $\exp(q_l)$. However, we need to find a careful combination of them for high performance in practice. All of these issues will be elaborated in the following sections.

B. Tension between Symmetric and Asymmetric Contention

(a) **Topological dependence.** A good combination of two CSMA operational parameters for high performance depends on contention topologies. O-DCF is designed to *autonomously* choose the combination of access probability and transmission length without explicit knowledge of topological information. Just for ease of explanation, we provide our description, assuming that flows are configured in either of the two “extreme” topologies: fully-connected (FC) for symmetric contention and flow-in-the-middle (FIM) for asymmetric contention (see Fig. 2), but O-DCF generally works well beyond these two topologies. Recall the two key design ideas of O-DCF: we first choose the initial access probability as a sigmoid function of the queue length and then let it experience BEB. To summarize, BEB is a key component in symmetric contention and the sigmoid function based access probability selection is crucial in asymmetric contention, and both are important in “mixture” topologies (all of which are presented below, respectively).

(b) **O-DCF: How exponential backoff helps.** In symmetric contention, the access probabilities among the contending flows should be reasonably low; otherwise, throughput will naturally degrade. Note that to guarantee fairness and high (long-term) throughput, a tiny p_l can work as it leads to almost no collision. This is because in that case a significantly long transmission length would recover the long-term throughput, as explained in oCSMA theory. However, such a combination will lead to a serious problem in short-term fairness, where a maximum bound on transmission length to guarantee short-term fairness is enforced in practice. For an automatic adaptation to contention level, we utilize BEB as a *fully distributed search process* for the *largest* access probability (i.e., the smallest CW size) that lets the links share the media efficiently in presence of collisions.

(c) **O-DCF: Why sigmoid function.** As opposed to symmetric contention, in asymmetric contention such as FIM-like topologies, almost no collision occurs and thus BEB rarely operates (we confirmed in Section V-B). More importantly, in this case, the starvation of the central flow is a major issue. To tackle this, we require that the CW size (or the access probability) of link l that solely contends with many other links should

be *small* (or *high*) and thus *prioritized* enough that the link l avoids rare channel access and even starvation. To provide such *access differentiation*, we focus on the fact that the flow in the middle, say l , typically has a longer queue than the outer flows. Let us denote the access probability p_l by some function of queue length q_l , i.e., $p_l = f(q_l)$. Thus, it is natural to design $f(q_l)$ to be *increasing* for access differentiation.

The question is what form of the function $f(q_l)$ is appropriate for high performance. To streamline the exposition, we proceed the discussion with the access probability rather than the CW. Toward efficient access differentiation, we start by the f 's requirements: for any link l and for $q_{\min} \leq q \leq q_{\max}$,

R1: $0 \leq f(q_l) \leq \bar{p} < 1$, where $f(q_{\min}) \approx 0$ and $f(q_{\max}) = \bar{p}$.

The largest access probability \bar{p} should be strictly less than one to prevent channel monopolization.

R2: $[f(q_l)]_2^4$ should span all the values in $\{1/2^i, i = 0, \dots, 9\}$ each of which corresponds to the CW sizes $\{2^{i+1} - 1, i = 0, \dots, 9\}$.

R3: The transmission lengths of the flows with heavy contention and those with light contention should be kept small not to exceed maximum aggregation size/time.

The requirement **R3** is important to prevent the central flow from being starved. In the FIM-like topologies, the flows experiencing heavy contention such as the central one in FIM has very rare chances to access the media. To guarantee (proportional) fairness, it is necessary for such flows to select long transmission lengths whenever holding the channel. However, as mentioned earlier, the transmission length should be bounded with the maximum aggregation size and/time for practical purposes such as short-term fairness. This implies that the central flow often needs to stop the transmissions before its required transmission length for optimal fairness is reached. Efficient flow differentiation, which prioritizes the flows with heavy contention in terms of access probability, significantly helps by reducing the required transmission length with a reasonable value (mostly shorter than the maximum transmission length) towards long-term fairness. An extreme case for asymmetric contention is IA scenario where two flows have asymmetric interference relationship. We will explain in Section IV-C that our flow differentiation helps a lot in providing (long-term) fairness in such a problematic scenario.

An intuitive way to realize flow differentiation is to set the access probability of link l to be $\exp(q_l)/K$ for some constant K . Then, the rule (5) enforces the transmission length to be around K , *irrespective of the contention levels of the flows* (i.e., **R3**). However, to satisfy **R1**, we use a slightly different function that has a sigmoidal form, $f(q_l) = \frac{\exp(q_l)}{\exp(q_l) + C}$, with some constant C . This function naturally makes the chosen access probability to be strictly less than one for any $q_{\min} \leq q \leq q_{\max}$, unlike $\exp(q_l)/K$. Clearly, the sigmoid function is not exponential over the entire $q_{\min} \leq q_l \leq q_{\max}$ values. However, it suffices to have an exponential form up to q_l' with $f(q_l') = 0.75$, since for a larger $q_l > q_l'$, the CW size approaches

⁴We denote by $[x]_2$ the $1/2^i$ for some integer i , which is closest to x , e.g., $[0.124]_2 = 1/2^3$ and $i = 3$.

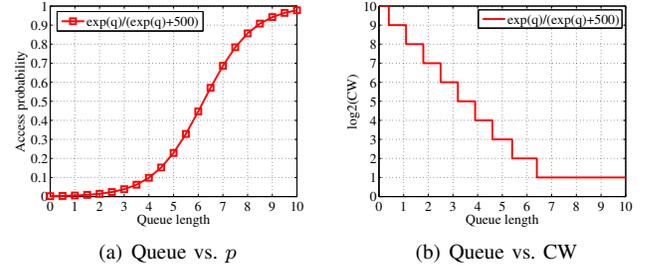


Fig. 3. Illustration of sigmoid function with respect to queue length.

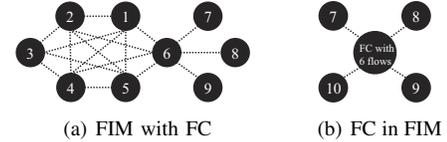


Fig. 4. Mixed topologies' conflicting graphs; (a) The flow 6 belongs to both FC and FIM topologies; (b) six flows within a FC group form a FIM topology with four outer flows.

one from the CW granularity (but, the access probability is set to be strictly less than one).

Then, the next question is the inflection point, determined by the constant C . We choose C around 500 due to the following reasons. First, for the resultant p to span the whole feasible values (as in **R2**), C should be greater than 500, i.e., $f(q_{\min}) = \frac{\exp(q_{\min})}{\exp(q_{\min}) + 500} \approx \frac{1}{512}$, where $q_{\min} = bQ_{\min} = 0.01$ (see Fig. 3(b) in terms of CW). Second, the parameter C determines the location of inflection point in the sigmoid function. For example, let us consider the FIM topology with four outer flows, where the central flow can have its queue length of up to five, while the outer flows have much smaller queue lengths usually less than two. To guarantee the exponential increase in that curve, the x -axis at the inflection point should be larger than five, implying that C larger than 500 is sufficient since $f(5) = \frac{\exp(5)}{\exp(5) + 500} < \frac{1}{2}$. However, too large C values directly result in too long transmission lengths from (5) because they may yield very low p 's.

(d) Mixture of contention levels. In practice, it is possible for a link to appear in a mixture of topologies with symmetric and asymmetric contention. We study this issue using the example scenarios in Fig. 4. First, in Fig. 4(a), flow 6 interferes with flows 1, \dots , 5 in a fully-connected fashion, and also with flows 7, 8, and 9 in a FIM-like fashion. Since it senses the transmission of both the remaining links forming the FC group and outer links forming the FIM group, its queue temporarily builds up, thus having a larger access probability than outer links due to the sigmoidal curve. However, it shares the medium equally with others in FC group so that their queue lengths increase together. Thus, in the worst case, BEB can prevent too aggressive access among the links within FC group. More importantly, *even the reduced access probability from BEB is kept larger than those of outer links* (see Section V-B for simulation results), thus still being sufficiently prioritized in the FIM topology, preserving proportional fairness. Similar trends

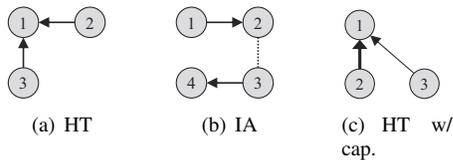


Fig. 5. Example topologies of imperfect sensing and packet capture. In network graphs, vertices represent nodes, dotted lines represent connectivity, and arrows represent flows. (a) HT; (b) IA; (c) HT with capture: a thin arrow for weak signal.

are also observed in Fig. 4(b).

(e) Aggressiveness control for session tails. Sessions dynamically enter and leave over time. As sessions' packets are served in a node (without new session arrivals), the CQ size naturally decreases, resulting in progressively less aggressive media access. In O-DCF, as the CQ size becomes zero, we record the MAQ size at that time, and access the media with the recorded MAQ size (not actual MAQ size) until all the packets at MAQ are served. This effectively solves the problem that the session tails may be processed with too low aggressiveness.

C. Imperfect Sensing and Capture Effect

We have so far described O-DCF for the perfect sensing case. However, sensing is often imperfect, as illustrated in Fig. 5. RTS-CTS-like virtual sensing in O-DCF is not a complete solution, because just a naive choice of access probability and transmission length may generate a self-feeding loop where collisions increase queue lengths, which in turn leads to more aggressive access and thus heavier collisions, especially under the CW adaptation (defined in Section V-A) [22]. Our queue based initial CW with BEB substantially lessens such bad impacts. In *HT*, if the queue lengths of hidden nodes are large, BEB lets each node increase its CW, so that with small time cost, a transmission succeeds. This successful transmission generally decreases the queue lengths, preventing CW from being too small due to queue based initial CW selection. In *IA*, collisions are asymmetric. Suppose that at some time, advantaged and disadvantaged flows have small and large queue lengths, respectively. Then, the disadvantaged flow will have a smaller initial CW, thus leading to successful transmissions and simultaneously the advantaged flows will hear CTS signalling (responding to RTS from the disadvantaged flow) and stop their attempts. This helps a lot in providing fairness between two asymmetric flows. The packet capture effects can be handled by our method similarly to the *IA* scenario, i.e., interference asymmetry. For example, in *HT with capture*, similar behaviors to the *IA* case occur between weak and strong nodes.

D. O-DCF Implementation

(a) Queue structure. We implemented O-DCF as an overlay MAC on legacy 802.11 hardware using our C-based software platform, which requires protocol implementation on top of MadWiFi device driver [31]. Due to a limited memory size of legacy network interface card (NIC), we implement CQ and MAQ at the user space level. Note that the scheduling from

MAQ to IQ is not an actual packet transmission to the media. To minimize the temporal gap between the service from MAQ and the actual transmission, we reduce the buffer limit of IQ to one through a device driver modification.

(b) O-DCF scheduler and parameter control. A scheduler in O-DCF schedules the packets enqueued at MAQ to send them into the 802.11 hardware, as shown in Fig. 1. Whenever IQ in 802.11 becomes empty, the scheduler in O-DCF is notified by a system call such as `raw socket` function and determines the next packet to be enqueued into IQ, by comparing the lengths of multiple per-neighbor MAQs. Meanwhile, the scheduler maintains CSMA parameters, such as CW, AIFS, and NAV values, for each MAQ's HOL (Head-Of-Line) packet. To facilitate packet-by-packet parameter control, we piggyback such parameters into the header of HOL packet, so that the modified driver can interpret and set them in the TXQ descriptor of an outgoing packet for the actual transmission.

(c) Long data transmission. When packet aggregation is not supported in legacy 802.11 hardware, as in our real testbed, such as 802.11a/b/g, we take the following approach. The O-DCF scheduler assigns different arbitration inter-frame spaces (AIFSSs) and CWs for packets inside the specified transmission length. Since AIFS defines a default interval between packet transmissions and the smallest CW indicates the shortest back-off time, this provides a prioritization for *back-to-back* transmissions until the given transmission length expires. Further, we exploit the network allocation vector (NAV) option that includes the time during which neighbors remain silent irrespective of sensing. This guarantees that even interfering neighbors unable to sense (due to, e.g., channel fluctuations) do not prevent the transmission during the reserved transmission length.

V. PERFORMANCE EVALUATION

A. Setup

We use both NS-3 simulation and real experiments on a 16-node wireless testbed. The simulations are conducted for the basic topologies as well as the topologies that are hard to manually configure in the real experiments. We perform testbed experiments for the topologies randomly made in a building as well as the topologies that are more sensitive to the real environments such as packet capture. The details of our testbed specification will be described in Section V-C. For all scenarios, we repeat ten times (each lasting for 100 seconds) and measure the goodput at receiver, and the length of error bars in all plots represents standard deviation. The packet size is 1000 bytes. The parameters used in O-DCF are as follows: the queue scaling constant $b = 0.01$, the queue bounds are $Q_{\min} = 1$, $Q_{\max} = 1000$, and $V = 400$ (recall that the dequeue rate from CQ to MAQ is $V/q_i^M(t)$). We consider fully-backlogged flows in most scenarios, except for one plot that shows how O-DCF works for dynamic session arrival and departure. See Remark 5.1 for our parameter choices and other evaluation results.

We compare (i) 802.11 DCF, (ii) two versions of oCSMA in theory, and (iii) DiffQ [7]. For the standard oCSMA, we

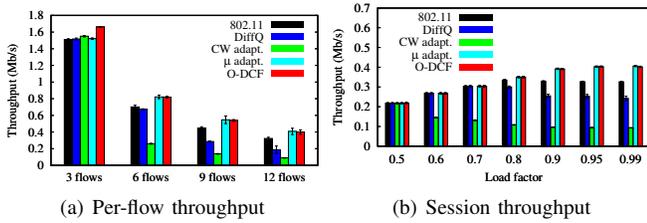


Fig. 6. (a) Average throughput in FC with 3, 6, 9, and 12 fully-backlogged flows (b) Session throughput for dynamic (thus, unsaturated) traffic.

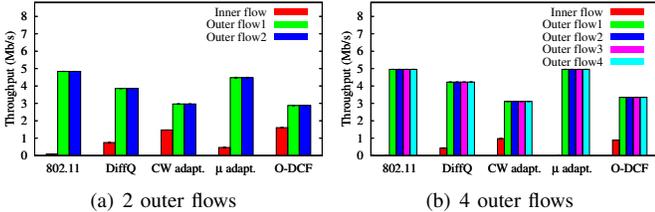


Fig. 7. FIM topology with two and four outer flows. In (a) and (b), the ideal proportional-fair throughput ratios of an outer flow and a middle flow are: 2:1 and 4:1, respectively.

test two versions (to show the effect of our automatic CSMA parameter combination in O-DCF: (a) *CW adaptation* in which we typically fix the transmission length μ with a single packet and control the access probability $p_i(t)$, such that $p_i(t) \times \mu = \exp(q_i(t))$ [1], and (b) *μ adaptation with BEB* (simply μ adaptation in this paper) in which we delegate the selection of $p_i(t)$ to 802.11 DCF and control $\mu_i(t) = \exp(q_i(t))/p_i(t)$. Note that to understand the effect of different methods for the adaptation of CWs, we evaluate μ adaptation with BEB using 802.11's CW size, and compare it with O-DCF. DiffQ is a *heuristic* queue-based MAC based on the 802.11e feature and schedules the interfering links with different priorities based on queue lengths.

B. Simulation Results: Basic Topologies

(a) Fully-connected: Impact of contention degrees. Fig. 6(a) summarizes the average per-flow throughput for the FC topologies (see Fig. 2(a)) for varying contention levels. The key point here is that O-DCF performs well *regardless of* the contention levels. All algorithms are good for the small number of flows. However, with increasing contention levels, the throughputs of DiffQ, 802.11, and CW adaptation decrease. μ adaptation performs similarly to O-DCF because in FC, they transmit multiple packets by enlarged transmission length after backing off with proper CW size adjusted by BEB. 802.11's bad performance in many-flow case is because the flows are unconditionally aggressive, where O-DCF's initial access probability is adapted over time.

(b) FIM: Impact of contention asymmetry. Next, we study the case of asymmetric contention, using the FIM topology (see Fig. 2(b)). Fig. 7 shows that O-DCF and CW adaptation dominate other algorithms. We observe that collisions are rare

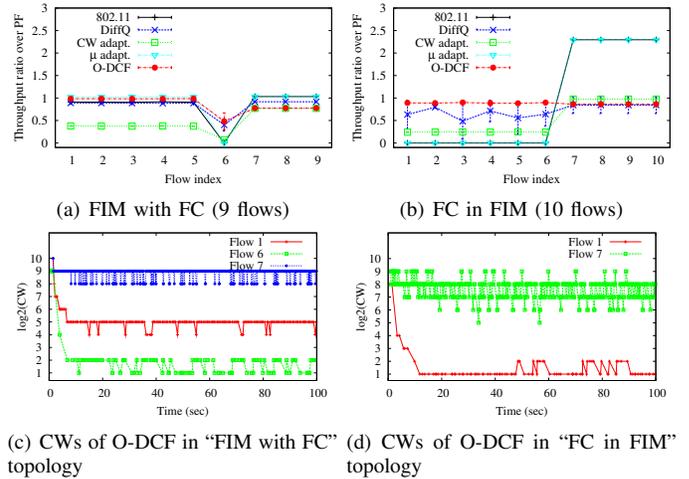


Fig. 8. Simulation. (a) - (b): Performance comparison in two mixed topologies; the flow indexes are the same as in Fig. 4; we plot the throughput ratio over optimality: the closer to one the ratio is, the better fairness is.;(c)-(d): CW traces of representative flows are plotted over time in O-DCF.

(less than $10\%^5$ in simulation, less than 2% in experiment), and thus BEB hardly happens. As reported in other papers [7], [9], the central flow experiences serious starvation in 802.11. In contrast to the symmetric contention as in FC, μ adaptation with BEB leads to the starvation of the central flow, since outer flows use 802.11's small CW to access the channel due to few collisions, which is too short to have enough overlapping of silent times and let the central flow transmit. Thus, in O-DCF, we can adapt the transmission length μ with a smarter CW selection method for high fairness. DiffQ resolves the starvation of the central flow well, but it shows suboptimal performance due to the heuristic setting of CW size.

(c) Mixed topology. We now consider the case when some flows belongs in part to symmetric and asymmetric contention (see Fig. 4). Fig. 8(a) shows the normalized per-flow throughput by the optimal PF share over the topology in Fig. 4(a). We observe that O-DCF outperforms others in terms of fairness. Note that the flow 6 placed between two different contention patterns experience some suboptimal fairness. This is because flow 6 should have low access probability due to the contention with flows 1-5, but it should also be prioritized over the flows 7-9. Our result shows that despite such conflicting situations, O-DCF tries to adjust the access probability and set the transmission length, so that the fairness is sustained to be reasonable, whereas other algorithms suffer from bad fairness. This is indeed achieved by our queue-length based initial CW selection with BEB in a distributed manner, which is verified by Fig. 8(c). Similar principles are applied to the topology in Fig. 4(b), whose (normalized) per-flow throughput and CW traces are shown in Fig. 8(b) and Fig. 8(d), respectively.

(d) Dynamic traffic. In addition, we evaluate the case with bursty traffic where active flow sessions can dynamically arrive

⁵This seems quite high, but the reason is that the total number of transmission attempts at central flow is small. Thus, the absolute number of collisions is small.

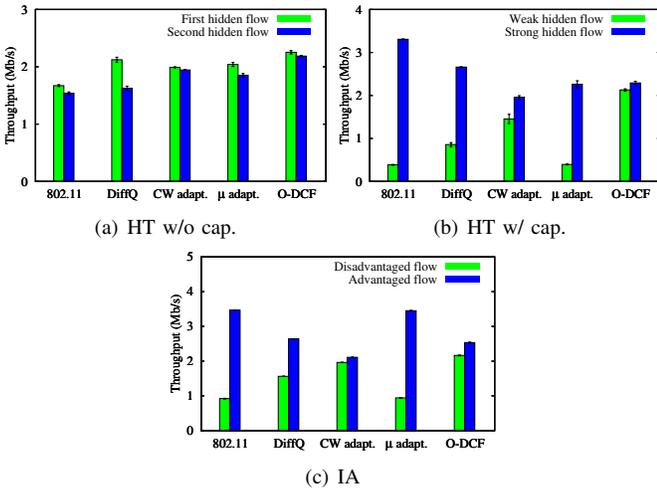


Fig. 9. Performance comparison in HT without packet capture (up and left) and with packet capture (up and right), and IA (down) topologies.

or leave the system. In a fully-connected topology with 12 flows, we generate multiple active sessions, each of them transferring a file of 1 MB size. The results are shown in Fig. 6(b), for different system load values. To obtain different load levels, we vary the session inter-arrival times, which is $1\text{MB}/(5.4\text{ Mb/s} \times \text{load factor}) \times 12$. Then, we evaluate system throughput vs. load as shown in the figure. We observe that, when the network load is low, all protocols can satisfy the flow traffic demands (see the case with load factor 0.5), and therefore all protocols have similar performance. However, as the load in the network increases, the system starts behaving in a more similar way to the fully-backlogged case, accumulating packets at the MAC layer queues as traffic demands overload system capacity. As an extreme case, for a network load of 0.99, we observe a very similar performance for all protocols to the one in Fig. 6(a) (i.e., the fully-backlogged case). In conclusion, with bursty traffic, low demands are easy to satisfy by any of the protocols, but with bursty traffic and higher demands, the benefits in the use of O-DCF become more evident, and are comparable to the ones already explained for the fully-backlogged case.

C. Experiments

In our 16-node testbed, each node is a netbook platform (1.66 GHz CPU and 1 GB RAM) running Linux kernel 2.6.31 and equipped with a single 802.11a/b/g NIC (Atheros chipset) running the modified MadWiFi driver for O-DCF's operation. To avoid external interference, we select a 5.805 GHz band in 802.11a. The default link capacity is fixed with 6 Mb/s, but we vary capacity or turn on auto rate adjustment, if needed.

(a) Imperfect sensing and capture effect. We investigate the impact of imperfect sensing in O-DCF with the topologies: HT, IA, and HT with capture, which are depicted in Fig. 5. As discussed in Section IV-C, we enable a virtual sensing, i.e., RTS/CTS signaling by default for better channel reservation before data transmission in all tested algorithms. Fig. 9 shows the

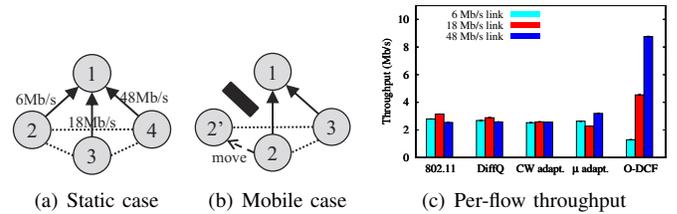


Fig. 10. Experiment. Single-hop network scenario consisting of clients and one AP; (a) static case: all nodes remain fixed with different PHY rates; (b) mobile case: node 2 moves away from node 1 at 60 seconds.

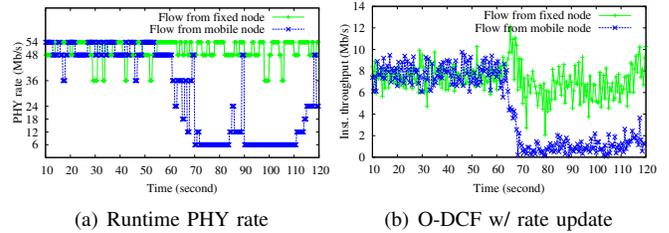


Fig. 11. Experiment. Performance evaluation in mobile scenario (Fig. 10(b)).

throughput results for the three topologies, respectively. First, in HT, O-DCF outperforms others in symmetric interference conditions by hidden nodes, as well as asymmetric conditions due to packet capture. Particularly, in O-DCF, collisions and thus BEB allow the hidden flows to access the media with larger CWs (i.e., less aggressiveness), resulting in many successful RTS/CTS exchanges. Fairness is guaranteed by the transmission length control. Second, in IA, our O-DCF shows very fair and high throughput, where the advantaged flow has a larger CW due to small backlog, and the disadvantaged flow also uses a larger CW due to BEB, responding well to collisions. As a result, both flows contend with sufficiently large CW values without heavy collisions, enjoying better RTS/CTS-based medium reservation. Similarly, we can ensure fairness in the scenario with packet capture. In μ adaptation, the advantaged flow accesses the channel using 802.11's small CW due to few collisions, which leads to a high probability of collisions and even starvation of the disadvantaged flow, and thus unfairness between two flows deepens. DiffQ cannot solve unfairness perfectly due to the still aggressive CW values. However, CW adaptation grants more aggressiveness to less-served flows, thus reducing the throughput gap between two flows in those scenarios, while still leaving a small gap from optimality.

(b) Heterogeneous channels. We consider two scenarios: (i) *static* in Fig. 10(a): nodes are stationary with different link rates, 6, 18, and 48 Mb/s, and (ii) *mobile* in Fig. 10(b): each node turns on the auto-rate functionality and two clients send their data to a single AP, and after 60 seconds, one of two clients (say node 2) moves away from AP (node 1). In adapting the rate, we employ the famous SampleRate algorithm [27] in MadWiFi driver. We measure the runtime PHY rate information updated every one second interval. First, in the static scenario, as shown in Fig. 10(c), we observe that only O-DCF can attain

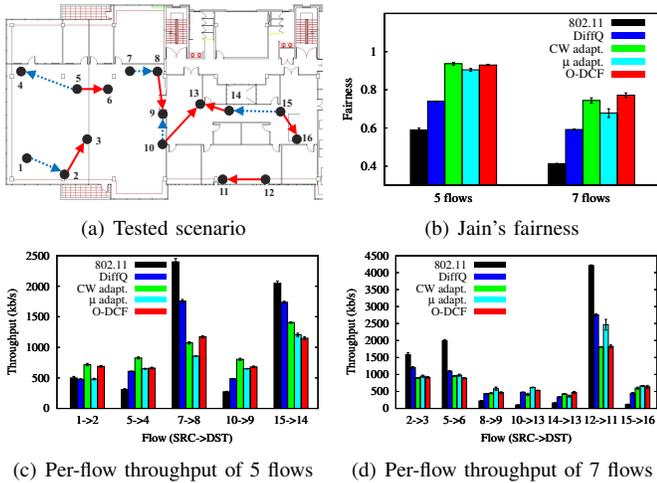


Fig. 12. Experiment. Tested topology and performance comparison; (a) 16 nodes denoted by triangles are distributed in the area of 40m x 20m; dotted (solid) arrows represent 5 (7) flows for the first (second) scenario. (b) Jain's fairness comparison. (c)-(d) Per-flow throughput distributions.

proportional fair rate allocation in an efficient manner, because of the consideration of link rate heterogeneity in the choice of CWs and transmission lengths, whereas the other protocols show severe inefficiency that the flows with higher rates are significantly penalized by that with lower rate. Besides the long-term fairness study in the static scenario, we examine the responsiveness of O-DCF to channel variations in the mobile scenario, shown in Fig. 11(a). We trace the instantaneous throughput of both fixed and mobile nodes in Fig. 11(b), where we see that the incorporation of the runtime PHY rate in O-DCF indeed helps in achieving throughput efficiency instantaneously even in the auto-rate enabled environment.

(c) Random topology experiment in a building. We now conduct a more general experiment using a 16-node testbed topology, as shown in Fig. 12(a), where we test two cases of five and seven concurrent flows. This random topology enables us to see how the algorithms perform in the mixture of hidden terminals and heavy contention scenarios including FIM scenario. Fig. 12(b) compares Jain's fairness achieved by all the algorithms for two scenarios. We find that over all the scenarios, O-DCF outperforms others in terms of fairness (up to 87.1% over 802.11 and 30.3% over DiffQ), while its sum utility is similar with others. The fairness gain can be manifested in the distribution of per-flow throughput, as shown in Fig. 12(c) and Fig. 12(d). O-DCF effectively prioritizes the flows with more contention degree (e.g., flow 10 \rightarrow 9 forms *flow-in-the-middle* with flows 7 \rightarrow 8 and 15 \rightarrow 14) and provides enough transmission chances to highly interfered flows (i.e., 8 \rightarrow 9, 10 \rightarrow 13, and 14 \rightarrow 13), compared to 802.11 DCF and DiffQ. The experimental topology is somewhat limited in size, tending to be full-connected. This leads to a small performance gap between oCSMA and O-DCF, but 802.11 DCF yields severe throughput disparities of more than 40 times between flows 12 \rightarrow 11 and 10 \rightarrow 13 in the second scenario. Compared to 802.11, DiffQ performs fairly well in the sense that it prioritizes

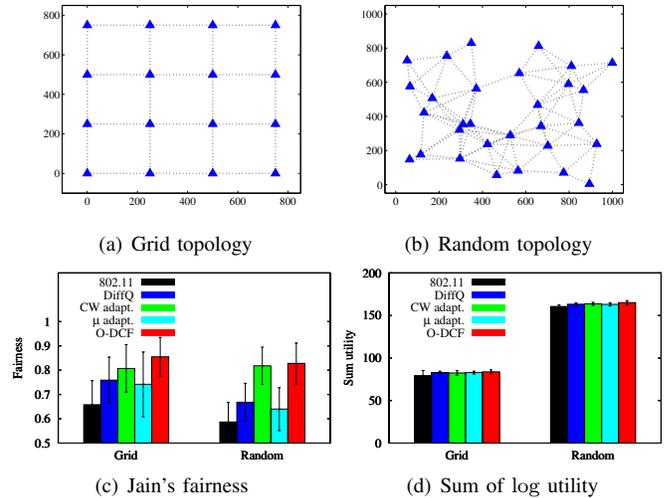


Fig. 13. Simulation. Large-scale topologies within 1000m \times 1000m area and each node denoted by triangles has transmission and sensing ranges equal to 280m; (a) 16 nodes form a grid network; (b) 30 nodes are deployed randomly; (c)-(d) fairness and sum utility comparison among tested algorithms.

highly interfered flows. However, its access prioritization is heuristic, so there is still room for improvement compared to O-DCF.

D. Simulation Results: Large Topologies

Finally, we extend the evaluation of O-DCF to two large-scale scenarios in a wider area. The first is a grid network where 16 nodes are apart from each other with 250m, and the second is a random network where 30 nodes are deployed randomly within an area of 1000m \times 1000m. These networks, depicted in Figs. 13(a) and 13(b), contain a mixture of the previously discussed problematic scenarios, such as HT, IA, and FIM, as well as highly interfering FC groups. For a given number of flows, we construct 10 scenarios with different random single-hop flows. In particular, we simulate 6 and 12 flows over grid and random networks, respectively.

Fig. 13(c) compares the Jain's fairness achieved by five algorithms. O-DCF enhances fairness up to 41.0% (resp., 29.9%) over 802.11 DCF and 24.0% (resp., 12.6%) over DiffQ in the random (resp., grid) network. Compared to the previous experiments, we observe more throughput disparities of contending flows under DiffQ and μ adaptation. In most scenarios, randomly chosen flows likely have problematic relationships such as HT and IA. This leads to lower fairness of DiffQ and μ adaptation than O-DCF and CW adaptation, since they are vulnerable to imperfect sensing scenarios, as examined in Section V-C(a). Moreover, O-DCF experiences much less collisions than CW adaptation, thus leading to better throughput. Finally, our O-DCF consistently shows a slightly better performance in terms of sum utility (see Fig. 13(d)).

Remark 5.1: First, the parameters of O-DCF are: b , Q_{\min} , Q_{\max} , and V . The value of b can be chosen by a typical step-size in many of optimization algorithms, say 0.01 or 0.001. The queue bound parameter is important, because $Q_l^M(t)$ is

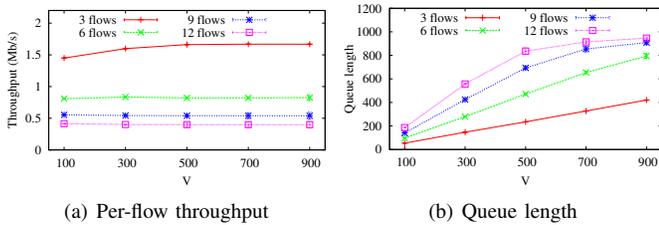


Fig. 14. Simulation. FC topology with 3, 6, 9 and 12 flows for different values of V .

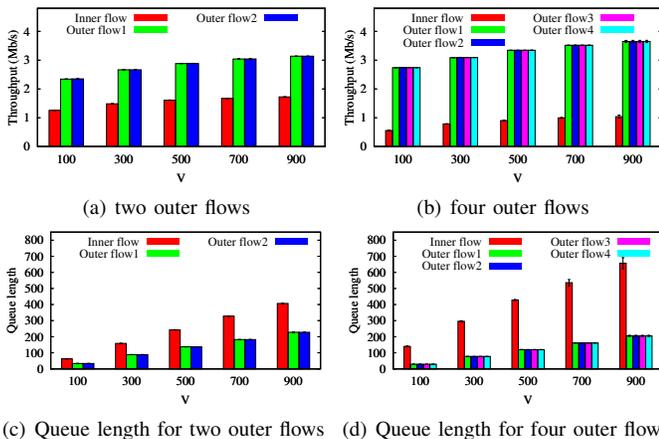


Fig. 15. Simulation. FIM topology with two and four outer flows for different values of V .

used to determine the injection rate from CQ to MAQ (i.e., $V/(bQ_i^M(t))$), connecting to the media aggressiveness. Thus, Q_{\min} and Q_{\max} are chosen just to avoid too small or large injection rates. Finally, it is intuitive that larger V leads to more sensitive response to the changes of $Q_i^M(t)$. This intuition has been proven in theory that V controls a tradeoff between long-term throughput and queueing delay [32]. We show through simulations such a tradeoff using FC and FIM topologies (see Fig. 14 and Fig. 15, respectively). In case of FC topology with 3 flows, we can observe that there is some gain in throughput with the increasing value V at the cost of increasing queueing delay. But, in case of FC topology with more flows (≥ 6) there is no gain in throughput even with larger V since wireless channel is already saturated by many contending flows. However, in case of FIM topology with two and four outer flows, larger V leads to throughput increment of all flows together with queueing delay increment since the wireless channel is not crowded with contending flows.

Second, it is not surprising that we observe that long transmissions in O-DCF sometimes increase the inter-TX delay of some flows. However, even in such a case, O-DCF still keeps substantial benefits in terms of long-term fairness. This confirms the results of prior work [3], [21], which has shown a tradeoff between short-term fairness and long-term efficiency in oCSMA algorithms. More importantly, the performance benefits in O-DCF can be found in critical scenarios where other CSMA protocols fail to achieve fairness, as shown in the

previous sections.

VI. CONCLUDING REMARKS

The major design issues to improve 802.11 DCF are contention window selection and transmission length control against network contention, imperfect sensing, channel heterogeneity, and packet capture without any message passing. We proposed a protocol, called O-DCF, inspired by the recent theory on CSMA to tackle these issues. However, we highlight that our work is not just an implementation of what has been developed in theory, covering many important engineering optimizations that have large impacts on the actual performance and deployment in practice. This paper shows that a combination of effective design ideas and implementation choices can actually let a theory-motivated protocol applicable in practice. We hope that our paper encourages and motivates other follow-up work to further bridge the gap between practice and theory.

REFERENCES

- [1] L. Jiang and J. Walrand, "A distributed CSMA algorithm for throughput and utility maximization in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 18, no. 3, pp. 960–972, Jun. 2010.
- [2] S. Rajagopalan, D. Shah, and J. Shin, "Network adiabatic theorem: An efficient randomized protocol for contention resolution," in *Proc. ACM SIGMETRICS*, 2009.
- [3] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor, "Towards utility-optimal random access without message passing," *Wirel. Commu. Mob. Comp.*, vol. 10, no. 1, pp. 115–128, Jan. 2010.
- [4] J. Ni, B. Tan, and R. Srikant, "Q-CSMA: Queue-length based CSMA/CA algorithms for achieving maximum throughput and low delay in wireless networks," in *Proc. IEEE INFOCOM*, 2010.
- [5] P. Marbach and A. Eryilmaz, "A backlog-based CSMA mechanism to achieve fairness and throughput-optimality in multihop wireless networks," in *Proc. Allerton Conf.*, 2008.
- [6] D. Skordoulis, Q. Ni, H.-H. Chen, A. P. Stephens, C. Liu, and A. Jamalipour, "IEEE 802.11n MAC Frame Aggregation Mechanisms for Next-Generation High-Throughput WLANs," *IEEE Wirel. Comm.*, vol. 15, no. 1, pp. 40–47, Feb. 2008.
- [7] A. Warrior, S. Janakiraman, S. Ha, and I. Rhee, "DiffQ: Practical differential backlog congestion control for wireless networks," in *Proc. IEEE INFOCOM*, 2009.
- [8] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang, "MACAW: A media access protocol for wireless LANs," in *Proc. ACM SIGCOMM*, 1994.
- [9] M. Garetto, T. Salonidis, and E. Knightly, "Modeling per-flow throughput and capturing starvation in CSMA multi-hop wireless networks," in *Proc. IEEE INFOCOM*, 2006.
- [10] X. Wang and K. Kar, "Throughput modelling and fairness issues in CSMA/CA based ad-hoc networks," in *Proc. IEEE INFOCOM*, 2005.
- [11] S.-J. Han, T. Nandagopal, Y. Bejerano, and H.-G. Choi, "Analysis of spatial unfairness in wireless LANs," in *Proc. IEEE INFOCOM*, 2009.
- [12] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *Proc. IEEE INFOCOM*, 2003.
- [13] F. Cali, M. Conti, and E. Gregori, "Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit," *IEEE/ACM Trans. Netw.*, vol. 8, no. 6, pp. 785–799, Dec. 2000.
- [14] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, "Idle Sense: an optimal access method for high throughput and fairness in rate diverse wireless LANs," in *Proc. ACM SIGCOMM*, 2005.
- [15] V. A. Siris and G. Stamatakis, "Optimal CWmin selection for achieving proportional fairness in multi-rate 802.11e WLANs: test-bed implementation and evaluation," in *Proc. ACM WinTECH*, 2006.
- [16] Y. Gruenberger, M. Heusse, F. Rousseau, and A. Duda, "Experience with an implementation of the idle sense wireless access method," in *Proc. ACM CoNEXT*, 2007.
- [17] Y. Kim, S. Choi, K. Jang, and H. Hwang, "Throughput enhancement of IEEE 802.11 WLAN via frame aggregation," in *Proc. VTC-Fall*, 2004.

- [18] W. Kim, H. K. Wright, and S. M. Nettles, "Improving the performance of multi-hop wireless networks using frame aggregation and broadcast for TCP ACKs," in *Proc. ACM CoNEXT*, 2008.
- [19] S. Sen, R. R. Choudhury, and S. Nelakuditi, "CSMA/CN: carrier sense multiple access with collision notification," in *Proc. ACM MobiCom*, 2010.
- [20] T. Li, M. K. Han, A. Bhartia, L. Qiu, E. Rozner, Y. Zhang, and B. Zarikoff, "CRMA: collision-resistant multiple access," in *Proc. ACM MobiCom*, 2011.
- [21] J. Lee, J. Lee, Y. Yi, S. Chong, A. Proutiere, and M. Chiang, "Implementing utility-optimal CSMA," in *Proc. Allerton Conf.*, 2009.
- [22] B. Nardelli, J. Lee, K. Lee, Y. Yi, S. Chong, E. Knightly, and M. Chiang, "Experimental evaluation of optimal CSMA," in *Proc. IEEE INFOCOM*, 2011.
- [23] A. Aziz, D. Starobinski, P. Thiran, and A. Fawal, "EZ-Flow: Removing turbulence in IEEE 802.11 wireless mesh networks without message passing," in *Proc. ACM CoNEXT*, 2009.
- [24] R. Laufer, T. Salonidis, H. Lundgren, and P. L. Guyadec, "XPRESS: A cross-layer backpressure architecture for wireless multi-hop networks," in *Proc. ACM MobiCom*, 2011.
- [25] A. Sharma and E. M. Belding, "FreeMAC: framework for multi-channel mac development on 802.11 hardware," in *Proc. ACM PRESTO*, 2008.
- [26] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma, "Performance of reliable transport protocol over IEEE 802.11 wireless LAN: Analysis and enhancement," in *Proc. IEEE INFOCOM*, 2002.
- [27] J. C. Bicket, "Bit-rate selection in wireless networks," Master's thesis, MIT, 2005.
- [28] G. Tan and J. Gutttag, "Time-based fairness improves performance in multi-rate wireless LANs," in *Proc. USENIX*, 2004.
- [29] A. Proutiere, Y. Yi, T. Lan, and M. Chiang, "Resource allocation over network dynamics without timescale separation," in *Proc. IEEE INFOCOM*, 2010.
- [30] Y. Yi and M. Chiang, *Next-Generation Internet Architectures and Protocols*. Cambridge University Press, 2011, chapter 9: Stochastic network utility maximization and wireless scheduling.
- [31] "Multiband Atheros Driver for WiFi," <http://madwifi-project.org/>.
- [32] M. J. Neely, "Super-fast delay tradeoffs for utility optimal fair scheduling in wireless networks," in *Proc. IEEE INFOCOM*, 2006.