A Game-Theoretic Study of CSMA/CA Under a Backoff Attack

Jerzy Konorski

***Abstract*—CSMA/CA, the contention mechanism of the IEEE 802.11 DCF medium access protocol, has recently been found vulnerable to selfish *backoff attacks* consisting in nonstandard configuration of the constituent backoff scheme. Such attacks can greatly increase a selfish station's bandwidth share at the expense of honest stations applying a standard configuration. The paper investigates the distribution of bandwidth among anonymous network stations, some of which are selfish. A station's obtained bandwidth share is regarded as a payoff in a noncooperative *CSMA/CA game*. Regardless of the IEEE 802.11 parameter setting, the payoff function is found similar to a multiplayer Prisoners' Dilemma; moreover, the number (though not the identities) of selfish stations can be inferred by observation of successful transmission attempts. Further, a *repeated CSMA/CA game* is defined, where a station can toggle between standard and nonstandard backoff configurations with a view of maximizing a long-term utility. It is argued that a desirable station strategy should yield a fair, Pareto efficient, and subgame perfect Nash equilibrium. One such strategy, called CRISP, is described and evaluated.**

***Index Terms*—Ad hoc LAN, game theory, MAC protocol, selfish behavior.**

I. Introduction

In a MOBILE ad hoc network (MANET), both user- and network-end functions are comprised in a mobile station, whose identity need not be known to other stations, and whose actions cannot be mandated by any other party. However, most MAC, routing and transport protocols have been designed with a fully cooperative user in mind, therefore offer little protection against noncooperative users. These circumstances create incentives for MANET stations to *misbehave*, i.e., depart from standard protocols, while guaranteeing them freedom from punishment. User misbehavior can be *selfish* (targeting an unfairly large share of network resources) or *malicious* (aimed at disrupting network operation, such as a denial-of-service attack). Orthogonally, one classifies misbehavior according to the affected layer of network architecture [12], [24].

This paper focuses on selfish MAC-layer misbehavior in a single MANET cell, an ad hoc wireless LAN (WLAN), employing the IEEE 802.11 DCF medium access protocol [13]. The key contention mechanism, CSMA/CA, features a distributed backoff scheme for collision avoidance and resolution. Selfish attacks against CSMA/CA aim at obtaining an unfairly large bandwidth share and can be launched by tampering with the network interface card (NIC) or by modifying NIC driver software. For example, a station may draw successive backoff durations from a different range than the standard prescribes. Such a *backoff attack* induces long-term unfairness on top of short-term one inherent in contention MAC protocols [16]; it is difficult to prevent, physically or administratively, as the local backoff scheme is entirely under a station's control. A possible approach is to provide disincentives to potential attackers. We view the stations' obtained bandwidth shares as payoffs in a noncooperative *CSMA/CA game* and let a station act as a player maximizing its payoff. The likely outcome is a *Nash equilibrium* (NE) [10], from which no player has an incentive to deviate. We analyze a one-shot and repeated CSMA/CA game, and propose a station strategy that makes backoff attacks by other stations ultimately yield lower payoffs.

Numerous studies of network protocols have taken an incentive-oriented game-theoretic approach, e.g., [1], [8], [19], [25], [30], [31]. It is especially fruitful in ad hoc networks, which lack central administration and serve anonymous (i.e., unaccountable) stations. Here, most studies have been devoted to eliciting cooperation in packet forwarding, e.g., [5], [23], [24], [34], [36]. Earlier work on MAC-layer misbehavior includes analyses of S-ALOHA [21], [22] and random token protocols [17], discussion of vulnerabilities of IEEE 802.11's physical and virtual carrier-sense mechanisms [4], and proposals of punishment- and reputation-based measures [20]. For hot-spot WLANs, [28] describes a trusted agent for statistical detection of misbehavior given genuine station identities. Close to our model is a recent game-theoretic study [7], which demonstrates how to enforce a Pareto efficient NE by means of a distributed penalization mechanism (cf. also [26]). Its implementation again requires genuine station identities, hard to come by in ad hoc networks, and on-the-fly jamming of selected frames sensed on the medium, a significant departure from the MAC standard. We believe that with anonymous stations and no trusted party, defense against selfish MAC-layer misbehavior is not possible with an IEEE 802.11 NIC applied "as is." However, we want to keep the MAC protocol intact.

The rest of this paper is organized as follows. Section II contains a brief reminder of the CSMA/CA contention mechanism, extends a well-known approximate model of an ad hoc CSMA/CA WLAN under saturation load, and discusses its solution under a backoff attack. In Section III, we describe a one-shot CSMA/CA game. Next, in Section IV we describe a repeated CSMA/CA game and argue that a fair, Pareto efficient,

Manuscript received June 4, 2004; revised April 7, 2005, and September 21, 2005; approved by IEEE/ACM Transactions on Networking Editor A. Orda. This work was supported in part by the U.S. Air Force Office of Scientific Research under Grant FA8655-04-1-3074 and in part by the Ministry of Education and Science, Poland, under Grant 1599/T11/2005/29.

The author is with the Department of Information Systems, Faculty of ETI, Gdansk University of Technology, ul.Narutowicza 11/12, 80-952 Gdansk, Poland (e-mail: jekon@eti.pg.gda.pl).

Digital Object Identifier 10.1109/TNET.2006.886298

TABLE I
SUMMARY OF MAIN NOTATION

n	subscript indicating a station ($n = 1,\ldots,N$)
$w = \langle w_{\min}, L\rangle$	backoff configuration
$\boldsymbol{w} = (w_1,\ldots,w_N)$	configuration profile
t, c, s	transmission, collision, and success probability in a slot
T, S	total transmission and success probability
b	payoff (a station's bandwidth share)
W	set of feasible backoff configurations
g, s, h	subscripts related to (stations with) greedy, selfish, and honest backoff configurations
x, y	number of stations playing selfish and greedy
x^*	coarse profile observability threshold
k, r	superscripts counting CRISP stages and phases
U	utility function
σ	station strategy
M	CRISP threshold
p	probability of CRISP playing selfish or greedy

and subgame perfect NE [10] should be sought. In Section V, we analyze a strategy called CRISP that yields such a NE and thereby ensures asymptotically fair and efficient bandwidth use. Section VI concludes the paper and outlines directions for further research. Table I summarizes the notation used throughout the paper.

II. CSMA/CA Performance at Saturation

In this section, we outline the network operation and calculate a station's obtained bandwidth share using Bianchi's Markovian model. Next we extend the model to capture some general properties of CSMA/CA under backoff attack.

A. CSMA/CA Contention Under IEEE 802.11 DCF

A station waits until the medium has been idle for a predefined DIFS period and then sets the local *backoff counter* to a random integer from $[0, CW)$, where CW is the current *contention window*. Initially, CW is set to a minimum value $w_{\min}$. The backoff counter is decremented each time the medium is sensed idle for a predefined *slot* period, the countdown being frozen whenever the medium is sensed busy and resumed after it is sensed idle for another DIFS period. If *basic access* method is used, a DATA frame is transmitted when the backoff counter expires. If successful, it is replied to by a recipient's ACK frame and CW is reset to $w_{\min}$. Otherwise, i.e., if multiple DATA frames collide, each of their senders infers a collision (as no ACK follows), doubles CW, sets the backoff counter to a random integer from $[0, CW)$, and starts another countdown. The doubling of CW stops at a maximum value $w_{\max}$; further consecutive collisions beyond a *retry limit* cause a transmission abort (we neglect this feature in our analysis, cf. [27], [37]). A SIFS period, shorter than DIFS, is defined to guarantee uninterrupted DATA + ACK exchange. In *RTS/CTS access* method, a station whose backoff counter has reached zero transmits an RTS frame; if successful, it is replied to by a recipient's CTS frame and a DATA frame transmission follows. If multiple RTS frames collide, their senders infer a collision (as no CTS frame follows) and behave similarly as above. A SIFS period guarantees uninterrupted RTS + CTS + DATA + ACK exchange.

Let L be the maximum number of consecutive collisions upon which CW is doubled, i.e., $w_{\max} = w_{\min} \cdot 2^L$; the pair $w = \langle w_{\min}, L\rangle$ will be called the *backoff configuration* (or *configuration* for short) at a station. E.g., $w = \langle 16, 6\rangle$, $\langle 32, 5\rangle$, and $\langle 64, 4\rangle$ are configurations recommended for FHSS- or OFDM-based, DSSS-based, and infrared-based PHY layers, respectively.

B. Stochastic Performance Model

Consider an ad hoc, single-hop WLAN with N stations using the IEEE 802.11 DCF protocol in the MAC layer. The single-hop assumption aims to factor out hidden stations and issues related to frame forwarding. It does not rule out RTS/CTS access, provided that DATA frames are long enough. Further assume that the network operates under saturation load, i.e., upon a successful transmission of a DATA frame, a station is immediately ready to transmit another one. The N stations can be thought of as engaged in large file transfers or transmission of CBR traffic. Note that selfish behavior is pointless under light or moderate load, where each station mostly gets all the bandwidth it requires. To set the stage for backoff attacks, we allow a transmitting station to remain anonymous to nonrecipients, its location obscured by mobility and inaccuracy of tracking devices, and its logical identity obscured by a fictitious MAC address.

As no exact model of a saturated CSMA/CA network exists, Bianchi's approximate model [6] serves as a first-order approach (for applications and refinements, see, e.g., [7], [27], [37], [39]). At the core lies an "independence hypothesis": a station with $w = \langle w_{\min}, L\rangle$ is represented by a Markov chain, assuming a constant collision probability c (i.e., the probability of another station's transmission attempt in a slot). According to [6], the steady-state probability of the station attempting a frame transmission in a slot is then $t = \left((w_{\min}+1)/2 + (w_{\min}/4)\cdot\sum_{l=1}^{L}(2c)^l\right)^{-1}$. Backoff freezing can be accounted for similarly as in [39], yielding

$$t = \frac{1-c}{1-c+\frac{w_{\min}-1}{2}+\frac{w_{\min}}{4}\cdot\sum_{l=1}^{L}(2c)^l}. \quad (1)$$

The t thus expressed as a function of c reflects back upon c via the "independence hypothesis," which leads to $c = 1-(1-t)^{N-1}$, a nonlinear equation in c.

Suppose now that each station n configures its backoff scheme individually with $w_n = \langle w_{n,\min}, L_n\rangle$. Though feasible configurations are only constrained by $w_{n,\min} \geq 1$ and $L_n \geq 0$, we assume throughout this section that $w_{n,\min} \geq 2$ ($w_{n,\min} = 1$ will be discussed separately). Station n is "more selfish" than station m if $w_{n,\min} \leq w_{m,\min}$ and $L_n \leq L_m$, at least one of the inequalities being strict. The relevant probabilities at station n, t_n and c_n, are bound by a relationship analogous to (1). To facilitate calculation note that by the "independence hypothesis," the probability of at least one station attempting a frame transmission in the present slot is

$$T = 1 - \prod_{n=1}^{N}(1-t_n). \quad (2)$$

Since $c_n = 1 - \prod_{m\neq n}(1-t_m)$, we have

$$t_n = \frac{T-c_n}{1-c_n}. \quad (3)$$

Substituting (3) into (1) yields for $n = 1, \dots, N$

$$\frac{1 + c_n \cdot (T-2)}{T - c_n} = \frac{w_{n,\min} + 1}{2} + \frac{w_{n,\min}}{4} \cdot \sum_{l=1}^{L_n} (2c_n)^l. \quad (4)$$

C. Bandwidth Shares

The c_n and t_n solving the nonlinear equations (2) through (4) determine station n's bandwidth share:

$$b_n = \frac{\tau_{\text{DATA\_payload}} \cdot s_n}{\tau_{\text{DIFS+RTS}} - \tau_{\text{slot}} + \tau_{\text{slot}}/T + \tau_X \cdot S} \quad (5)$$

where $s_n = t_n \cdot (1 - c_n)/T$ is station n's *success probability* (probability of a successful frame transmission in a nonempty slot), $S = \sum_{m=1}^{N} s_m$ is the total success probability, $\tau_{(\cdot)}$ is a specified time or frame transmission duration, and $X = \text{SIFS} + \text{CTS} + \text{SIFS} + \text{DATA} + \text{SIFS} + \text{ACK}$. For basic access, substitute DATA for RTS and take $X = \text{SIFS} + \text{ACK}$. Note that in existing IEEE 802.11 settings, $\tau_{\text{RTS+DIFS}} > \tau_{\text{slot}}$ and $\tau_{\text{DATA+DIFS}} > \tau_{\text{slot}}$.

Let $\boldsymbol{w} = (w_1, \dots, w_N)$ be a network-wide backoff *configuration profile*. If multistability of CSMA/CA (and unwelcome short-term unfairness it implies [27]) is to be avoided then (2)–(4) should have a unique solution $(t_n^o(\boldsymbol{w}), n = 1, \dots, N)$. A sufficient condition for this, as well as partial characterization of the solution, is given in the following proposition whose full proof appears in [18].

Proposition 1: Let $\tilde{n}$ be a station with the smallest $w_{n,\min}$ among all the stations and the smallest L_n among the stations with the smallest $w_{n,\min}$; let $w_{\tilde{n}} = \langle \widetilde{w}_{\min}, \widetilde{L} \rangle$. If

$$\widetilde{w}_{\min} \geq 1 + \sqrt{2 \cdot \widetilde{w}_{\min} \cdot \mathbf{1}_{\widetilde{L}>0}} \quad (6)$$

where $\mathbf{1}_{(\cdot)}$ is the indicator function, then the following holds:

(i) The solution $(t_n^o(\boldsymbol{w}), n = 1, \dots, N)$ exists and is unique.

(ii) Let $w' < w$, $\boldsymbol{w} = (w_1, \dots, w_{n-1}, w, w_{n+1}, \dots, w_N)$, and $\boldsymbol{w}' = (w_1, \dots, w_{n-1}, w', w_{n+1}, \dots, w_N)$. Then for the resulting bandwidth shares, $b_n(\boldsymbol{w}') > b_n(\boldsymbol{w})$.

Sketch of proof: Tedious calculation shows that if (6) holds then for $T \geq \widetilde{T} = 2/(\widetilde{w}_{\min} + 1)$ each of the equations (4) has a unique root $c_n^*(T, w_n)$ in $[0, T)$, while for $T < \widetilde{T}$ some of them have none; moreover, the corresponding $t_n^*(T, w_n)$ is continuous and nonincreasing in T. Thus, (2) becomes $T = G(T, \boldsymbol{w})$, where $G(T, \boldsymbol{w}) = 1 - \prod_{n=1}^{N}(1 - t_n^*(T, w_n))$ is also continuous and nonincreasing in T. By inspection, $t_n^*(\widetilde{T}, w_n) = \widetilde{T}$ if $w_{n,\min} = \widetilde{w}_{\min}$, therefore $G(\widetilde{T}, \boldsymbol{w}) \geq \widetilde{T}$. Hence, the root $T^o(\boldsymbol{w})$ of (2) fulfills $T^o(\boldsymbol{w}) \geq \widetilde{T}$, inducing a unique $t_n^o(\boldsymbol{w}) = t_n^*(T^o(\boldsymbol{w}), w_n)$. This yields part (i). For part (ii), write

$$b_n(\boldsymbol{w}) = \left(\alpha_1 + \frac{\alpha_1 \sum_{m \neq n} s_m(\boldsymbol{w}) + \alpha_2 + \alpha_3/T^o(\boldsymbol{w})}{s_n(\boldsymbol{w})}\right)^{-1} \quad (7)$$

where α_1, α_2, and α_3 are positive constants determined from (5), and $s_m(\boldsymbol{w})$ is station m's success probability under configuration profile $\boldsymbol{w}$. Careful analysis of (3) and (4) shows that $T^o(\boldsymbol{w}) < T^o(\boldsymbol{w}')$, $s_n(\boldsymbol{w}') > s_n(\boldsymbol{w})$, and $s_m(\boldsymbol{w}') < s_m(\boldsymbol{w})$ for $m \neq n$. Thus, the fraction in (7) decreases when $\boldsymbol{w}$ is replaced by $\boldsymbol{w}'$. □

Condition (6) excludes only configuration profiles with $\widetilde{w}_{\min} \leq 3$ and $\widetilde{L} > 0$, e.g., the presence of a station n with $w_n = \langle 2, 1 \rangle$ is not allowed unless there is another station m with $w_m = \langle 2, 0 \rangle$. Fortunately, even with the threat of multistability put aside, such anomalous profiles do not present practical backoff attack scenarios: apparently, configuring as selfishly as $w_{n,\min} \leq 3$ station n is determined to grab as large a bandwidth share as possible and would not want to diminish it by configuring $L_n > 0$ i.e., agreeing to double CW upon collisions. Part (ii) addresses incentives: since configuring selfishly trades increased collision rate for increased transmission opportunity, it would seem that the net effect depends on the configuration at the other stations; yet a backoff attack incentive turns out to be always present.

III. CSMA/CA Game

When selecting $w_n = \langle w_{n,\min}, L_n \rangle$ so as to maximize b_n, station n has no prior knowledge of the other stations' selections; yet b_n depends on them, as well as on w_n. An N-player game thus arises in which bandwidth shares are viewed as players' payoffs. It is reasonable to assume that the game is noncooperative in that no binding agreements can be reached between the players as to their future play.

A. Game-Theoretic Framework

We begin with a few game-theoretic definitions [10], retaining the previous notation and terminology.

Definition 1: A *game* is a triple $(\{1, \dots, N\}, W, b)$, where $\{1, \dots, N\}$ is the set of players (stations), W is the set of feasible actions (backoff configurations), and b: $W^N \to \mathbf{R}^N$ is a payoff function. Each station n selects a backoff configuration $w_n \in W$ and subsequently receives a payoff (bandwidth share) $b_n(\boldsymbol{w})$ dependent on the configuration profile $\boldsymbol{w} = (w_1, \dots, w_N)$.

Definition 2: A configuration w is *dominant* for station n if $b_n(w_1, \dots, w_{n-1}, w, w_{n+1}, \dots, w_N) \geq b_n(w_1, \dots w_{n-1}, w', w_{n+1}, \dots, w_N)$ for all $w_1, \dots, w_{n-1}, w', w_{n+1}, \dots, w_N \in W$ with $w' \neq w$. It is *strictly dominant* if all the inequalities are strict.

Definition 3: A *Nash equilibrium* (NE) is a configuration profile $\boldsymbol{w} = (w_1, \dots, w_N)$ at which $b_n(\boldsymbol{w}) \geq b_n(w_1, \dots, w_{n-1}, w', w_{n+1}, \dots, w_N)$ for all $n = 1, \dots, N$ and $w' \in W$ with $w' \neq w_n$. A NE is *strict* if all these inequalities are strict.

At a NE, each station selects a best reply to the other stations' selected configurations, a likely outcome if all the players are rational (i.e., only interested in maximizing own payoffs) and their rationality is common knowledge [10], [14]. If a configuration w is strictly dominant for each station then the configuration profile $(w, \dots, w)$ (denoted all-w) is a unique and strict NE. If it is nonstrictly dominant, the NE need not be unique or strict.

Definition 4: A configuration profile $\boldsymbol{w} \in W^N$ is *Pareto inferior* to another configuration profile $\boldsymbol{w}' \in W^N$, if $b_1(\boldsymbol{w}) \leq b_1(\boldsymbol{w}'), \dots, b_N(\boldsymbol{w}) \leq b_N(\boldsymbol{w}')$ with at least one strict inequality. It is *Pareto efficient* if it is not Pareto inferior to any other configuration profile, and *fair* if $b_1(\boldsymbol{w}) = \dots = b_N(\boldsymbol{w})$.

From a global performance viewpoint, a fair and Pareto efficient configuration profile is a desirable outcome as a form of "cooperative equilibrium." Unfortunately, it does not need to coincide with a NE.

B. Greedy, Selfish, and Honest Configurations

Let $W = \{w = \langle w_{\min}, L\rangle | 1 \leq w_{\min} \leq \overline{w}_{\min}, 0 \leq L \leq \overline{L}\}$ and let $w_{\rm h} = \langle \overline{w}_{\min}, \overline{L}\rangle$ be a configuration prescribed by the IEEE 802.11 standard. A station is said to play *honest* if it sticks to $w_{\rm h}$, and *greedy* if it selects $w_{\rm g} = \langle 1, 0\rangle$; the latter configuration, so far excluded from considerations, has the backoff scheme disengaged (as permitted by some existing NIC driver software). Selecting $w_{\rm g}$ or any other configuration in the presence of another station playing greedy yields a zero payoff; however, selecting $w_{\rm g}$ with no other station playing greedy yields a 100% success probability and the highest possible payoff, further denoted $b_{\rm g}$ [for RTS/CTS access, $b_{\rm g} = \tau_{\rm DATA\_payload}/\tau_{\rm DIFS+RTS+X}$, whereas for basic access, $b_{\rm g} = \tau_{\rm DATA\_payload}/\tau_{\rm DIFS+DATA+X}$, cf. the notation following (5)]. Therefore, $w_{\rm g}$ is nonstrictly dominant for each station. Any configuration profile with at least one station selecting $w_{\rm g}$ is a nonstrict NE.

Given that each of these Nash equilibria is either Pareto inefficient or unfair, and that most of them yield zero payoffs, one may ask whether there are other configurations the players might be interested in selecting. Consider a *restricted game* with the constraint $w_{\min} > 1$, i.e., no station plays greedy. The payoffs can now be calculated as in Section II, provided that (6) holds. A station selecting $w_{\rm s} = \langle 2, 0\rangle$ is said to play *selfish*. By part (ii) of Proposition 1, $w_{\rm s}$ is strictly dominant for each station, so that the configuration profile all-$w_{\rm s}$ is a unique and strict NE. Unlike $w_{\rm g}$, $w_{\rm s}$ can be selected without fear that any other station responding in kind will reduce the received payoff to zero: as demonstrated below, two or more stations playing selfish can still count on nonzero payoffs. We conclude that each station has reasons to either play greedy or selfish (if it has the desire and ability to self-configure the backoff scheme), or honest (if it has no such desire or ability). In the former case it plays greedy when no station is expected to respond in kind, and selfish otherwise.

Definition 5: A *CSMA/CA game* conforms to Definition 1 with $W = \{w_{\rm g}, w_{\rm s}, w_{\rm h}\}$.

To further characterize the restricted CSMA/CA game, consider a practical scenario where x out of the N stations play selfish, whereas the other $N - x$ play honest. Hence, the desirable outcome is all-$w_{\rm h}$. Let $b_{\rm h}(N, x)$ and $b_{\rm s}(N, x)$ be an honest and selfish station's bandwidth share, respectively, with an analogous notation $s_{\rm h}(N, x)$, $s_{\rm s}(N, x)$, and $T(N, x)$. The $s_{\rm h}(N, x)$ and $s_{\rm s}(N, x)$ are only configuration dependent and can be obtained from (2)through (4). Table II lists sample values for $N = 5$, 10, 15, and 20 along with the total success probability $S(N, x) = (N - x) \cdot s_{\rm h}(N, x) + x \cdot s_{\rm s}(N, x)$, assuming $w_{\rm h} = \langle 16, 6\rangle$. These values uniquely determine $T(N, x)$ since in the equation defining s_n, both t_n and $1 - c_n$ are nonincreasing in T [18], hence the right-hand side is monotonously decreasing in T.

Turning to bandwidth shares, we find that $b_{\rm s}(N, N) < b_{\rm h}(N, 0)$ for any existing IEEE 802.11 setting. Indeed, from (5) a sufficient condition for this is

TABLE II
SUCCESS PROBABILITIES WITH $w_s = \langle 2, 0\rangle$ AND $w_{\rm h} = \langle 16, 6\rangle$

	$s_{\rm s}(N, x)$ and $s_{\rm h}(N, x)$ (%)							
x	$N = 5$		$N = 10$		$N = 15$		$N = 20$	
0		17.43		8.14		5.21		3.80
1	96.95	0.27	93.37	0.27	89.95	0.27	86.67	0.27
2	32.96	0.13	32.34	0.13	31.75	0.13	31.16	0.13
3	17.82	0.10	17.57	0.10	17.32	0.10	17.08	0.10
4	11.60	0.07	11.41	0.07	11.23	0.07	11.05	0.07
5	8.31		8.16	0.07	8.03	0.07	7.88	0.07
10			2.92		2.85	0.07	2.78	0.07
15					1.56		1.52	0.07
20							1.00	

	$S(N, x)$ (%)			
x	$N = 5$	$N = 10$	$N = 15$	$N = 20$
0	87.13	81.41	78.15	75.93
1	98.04	95.79	93.71	91.74
2	66.33	65.76	65.25	64.73
3	53.66	53.43	53.19	52.95
4	46.45	46.09	45.72	45.36
5	41.54	41.14	40.86	40.46
10		29.16	28.79	28.44
15			23.41	23.06
20				20.07

$$\kappa > 1 + \frac{s_{\rm s}(N, N)/T(N, 0) - s_{\rm h}(N, 0)/T(N, N)}{s_{\rm h}(N, 0) - s_{\rm s}(N, N)} \quad (8)$$

where $\kappa = \tau_{\rm DATA+DIFS}/\tau_{\rm slot}$ for basic access and $\kappa = \tau_{\rm RTS+DIFS}/\tau_{\rm slot}$ for RTS/CTS access. With $w_{\rm h} = \langle 16, 6\rangle$ and $N > 2$, the right-hand side of (8) remains below five, whereas κ ranges from over six (in a 54 Mb/s IEEE 802.11a setting with RTS/CTS access, where $\kappa \approx 6.33$) to over 200 (in a 1-Mb/s IEEE 802.11 setting with basic access and 1500-byte DATA frames, where $\kappa \approx 240$). With $w_{\rm h} = \langle 32, 6\rangle$ or $\langle 64, 6\rangle$ the conclusion remains the same. Hence, the unique NE all-$w_{\rm s}$ is Pareto inferior to all-$w_{\rm h}$, the only fair and Pareto-efficient configuration profile. This creates exactly the same conflict as does the well-known Prisoners' Dilemma [10], [21]: self-interest dictates that each station select $w_{\rm s}$, yet all the stations would be better off selecting $w_{\rm h}$. Calculating similarly as above we find that $b_{\rm s}(N, x)$ and $b_{\rm h}(N, x)$ decrease in x and N in existing IEEE 802.11 settings, which fits in with the definition of a multiplayer Prisoners' Dilemma [38].[1]

A look at Table II reveals two IEEE 802.11 setting independent properties of the restricted game, collectively referred to as *coarse profile observability*. First, $S(N, x)$ varies visibly with x, but little with N. One can define a threshold x^* such that for all N, N', and $x \leq x^*$, the relative difference between $S(N, x)$ and $S(N', x + 1)$ is significant; e.g., if 10% difference is considered significant then $x^* = 4$. Second, if N is not too large (e.g., not exceeding 20, a reasonable WLAN size) then $s_{\rm h}(N, 0)$ is distinctly positive, while $s_{\rm h}(N, x) \approx 0$ for $x > 0$.

Coarse profile observability brings valuable information into the game: even though station n cannot guess the other stations' configurations, it can infer x with a certain granularity. If it plays selfish, observation of s_n permits to distinguish $x = 1, \ldots, x = x^*$, and $x > x^*$. If it plays honest, a similar distinction follows

[1]Note that the nonrestricted CSMA/CA games also bears some resemblance to a multiplayer Prisoner's Dilemma in that the only fair NE is Pareto inefficient.

from the observed total success probability S, whereas observation of s_n enables the distinction between $x = 0$ and $x > 0$. (Even with anonymous stations, a long uninterrupted transmission followed by a short one after a SIFS period reveals a successful DATA + ACK exchange, so that S obtains as the proportion of nonempty slots with successful transmission attempts.)

In the presence of greedy play, x cannot be observed. However, letting $y = |\{m|w_m = w_\mathrm{g}\}|$ and observing S, a station can distinguish $y = 0$, $y = 1$, and $y > 1$ as indicated by $0 < S < 100\%$, $S = 100\%$, and $S = 0$, respectively. We summarize the properties of the CSMA/CA game as follows:

- any configuration profile with $y > 0$ is a nonstrict NE, either unfair or Pareto inefficient;
- coarse profile observability permits each station to distinguish $y = 0, y = 1$, and $y > 1$, and if $y = 0$, to distinguish $x = 0, \ldots, x = x^*$, and $x > x^*$;
- in existing IEEE 802.11 settings, the restricted CSMA/CA game is a multiplayer Prisoners' Dilemma.

IV. Repeated CSMA/CA Game

Although the Nash equilibria of the CSMA/CA game are either unfair or Pareto inefficient, game theory promises a more satisfactory outcome for a related repeated game. Let our game proceed in *stages* of fixed length. A stage should last long enough to approach the steady state values (5); simulations similar to those in Section V-C show that this happens after a few hundred to a few thousand CSMA/CA contentions. For each stage $k = 1, 2, \ldots$, station n selects a configuration $w_n^k \in W = \{w_\mathrm{g}, w_\mathrm{s}, w_\mathrm{h}\}$, which it maintains throughout the stage.[2] Thus, $\boldsymbol{w}^k = (w_1^k, \ldots, w_N^k)$ is the configuration profile in stage k and $(\boldsymbol{w}^1, \ldots, \boldsymbol{w}^k)$ is the *play path* up to stage k. An *observable* play path $\boldsymbol{\pi}$ only contains those characteristics of configuration profiles that are public knowledge, as they can be inferred via coarse profile observability. That is, $\boldsymbol{\pi} = (x^1, \ldots, x^k, \ldots)$, where for convenience we adopt the following notation, letting $y^k = |\{m|w_m^k = w_\mathrm{g}\}|$:

$$x^k = \begin{cases} \left|\{m|w_m^k = w_\mathrm{s}\}\right|, & \text{if } y^k = 0 \text{ and } x \le x^*, \\ > x^*, & \text{if } y^k = 0 \text{ and } x > x^*, \\ > N, & \text{if } y^k = 1, \\ \infty, & \text{if } y^k > 1. \end{cases} \tag{9}$$

Similarly as in [10], [14] we define a *strategy* σ_n as a function $\sigma_n \colon \Pi \to \Delta_W$ that specifies the probability distribution of selected w_n^k given an observed play path $\boldsymbol{\pi}$ up to stage $k-1$, where Π is the set of all observable play paths of finite length and Δ_W is the set of all probability distributions over W. The obtained *stage payoffs* are

$$b_n^k = \begin{cases} b_\mathrm{h}(N, x^k), & \text{if } x^k < N \text{ and } w_n^k = w_\mathrm{h}, \\ b_\mathrm{s}(N, x^k), & \text{if } 0 < x^k \le N \text{ and } w_n^k = w_\mathrm{s}, \\ b_\mathrm{g}, & \text{if } x^k => N \text{ and } w_n^k = w_\mathrm{g}, \\ 0, & \text{if } x^k = \infty \text{ or } \left(x^k => N \text{ and } w_n^k \neq w_\mathrm{g}\right). \end{cases} \tag{10}$$

[2]Throughout the paper, superscripts count time periods; for power exponents we use the notation $(a)^b$.

where b_g was defined in Section III-B. Any initial play path $\boldsymbol{\pi} \in \Pi$ determines past stage payoffs through (10) and, jointly with the strategy profile $(\sigma_1, \ldots, \sigma_N)$, induces a probability distribution μ_π of future stage payoffs. If all stage payoffs are equally weighted (there is no discounting of future payoffs), station n's long-term satisfaction, or *utility*, can be quantified by the following liminf-type asymptotic [15]:

$$U_n(\sigma_1, \ldots, \sigma_N|\pi) = \lim_{k\to\infty} \inf \left\{E_{\mu_\pi} u_n^k, E_{\mu_\pi} u_n^{k+1}, \ldots\right\} \tag{11}$$

where $u_n^k = (b_n^1 + \ldots + b_n^k)/k$. Since the u_n^k are bounded, so are their expectations, hence the limit exists and lower bounds the long-term per stage payoff average. Note that the inferior operator can be dropped if the u_n^k converge in probability to a constant as k increases.

Ideally, each station plays a strategy that, from any stage onwards, is the best reply to the other stations' strategies in the sense of (11) regardless of the play path so far, and the resulting stage payoffs asymptote to $b_\mathrm{h}(N, 0)$, i.e., ultimately all-w_h persists. This is captured by the notions of fairness, Pareto efficiency, and subgame perfection.

Definition 6: A strategy profile $(\sigma_1, \ldots, \sigma_N)$ is *fair and Pareto efficient* if

$$U_n(\sigma_1, \ldots, \sigma_N|\pi) = b_\mathrm{h}(N, 0) \tag{12a}$$

and a *subgame perfect NE* if

$$U_n(\sigma_1, \ldots, \sigma_N|\pi) \ge U_n(\sigma_1, \ldots, \sigma_{n-1}, \sigma, \sigma_{n+1}, \ldots, \sigma_N|\pi) \tag{12b}$$

for all $n = 1, \ldots, N$, $\boldsymbol{\pi} \in \Pi$, and any strategy σ of station n.

We seek a strategy σ^* such that the strategy profile $(\sigma^*, \ldots, \sigma^*)$ (also denoted all-σ^*) fulfills (12). Moreover, σ^* should only depend on observable play paths. From a practical viewpoint, σ^* should be simple (e.g., only depend on recent play), yet responsive to variable play of other stations. If found, σ^* can be made publicly known (note that the MAC standard remains unaffected, as toggling between w_h, w_s, and w_g is performed by user accessible software). This marks a change in perspective: instead of honest, selfish, and greedy stations we now speak of *standard* and *invader* stations. A standard station expects a fair bandwidth share and plays σ^* (not necessarily honest in each stage). An invader deviates from σ^* hoping for a more-than-fair bandwidth share. Condition (12a) ensures that while σ^* allows selfish or greedy play, it is by and large cooperative enough to achieve $b_\mathrm{h}(N, 0)$ in the absence of invaders. Condition (12b) discourages deviations from σ^*; yet they may occasionally occur due to a station's lapse of good judgment. In such cases, condition (12a) creates a friendly "learn and reform" environment: when an invader learns that it loses by deviating from σ^* and so reverts to σ^*, its stage payoffs asymptote to $b_\mathrm{h}(N, 0)$, as do the other stations'.

V. CRISP Strategy

Seeking a σ^* subject to (12) for the repeated restricted CSMA/CA game (with $W = \{w_\mathrm{s}, w_\mathrm{h}\}$) amounts to enforcing cooperation in the Prisoners' Dilemma, a classical game theory topic especially for the two-player case[3] [11], [29], [33], [35]. Most solutions draw on the idea of *generous tit-for-tat* [3], which applies a dose of generosity (playing honest regardless of other players' actions) in order to instill cooperation, while administering enough punishment for failure to cooperate. Note, however, that in the CSMA/CA game, station n has only imperfect information (x^k) on past configuration profiles; therefore it is not certain which other stations have played selfish, and even if it were, it would be unable to punish them selectively. The presence of the greedy configuration $w_{\rm g}$ further complicates the picture. Our candidate strategy σ^*, called CRISP (*Cooperation via Randomized Inclination to Selfish/Greedy Play*), inclines to selfish or greedy play in the presence of invaders, but in the absence thereof ultimately plays honest—thus imposes self-punishment in order to punish an invader; the latter can accept a diminished utility or revert to CRISP.

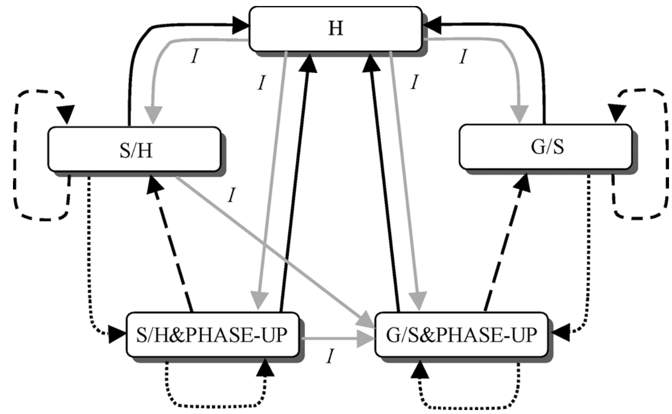

Fig. 1. CRISP state transition graph; solid, dashed, and dotted arrows correspond to downward trend, no trend, and upward trend in the play path, respectively; I marks invader-only transitions.

A. Strategy Description

A station adopting CRISP attempts to maintain all-$w_{\rm h}$ by selecting $w_{\rm h}$ if no greedy or selfish play by some other station has been detected in recent past. If selfish play has been detected, CRISP starts toggling between selfish and honest play, whereas if greedy play has been detected, CRISP toggles between greedy and selfish play. A public-knowledge infinite sequence ($p^r, r = 1, 2, \ldots$) with $p^r \in (0, 1)$, and a public-knowledge integer threshold M are defined; the latter controls the selfish/honest toggling ($1 \leq M \leq x^*$, where x^* is the threshold defined in Section III-B). The x^k, inferred by a station in stage k via coarse profile observability, are compared to 0 and M. Accordingly, only $x^k = 0$, $x^k = 1..M$, $x^k => M$, $x^k => N$, and $x^k = \infty$ are distinguished [cf. (9)]. Let this order be symbolized by "$\prec$" e.g., $> M \prec > N$.

The play proceeds in *phases*, each spanning a number of stages. We shall keep a uniform phase numbering, although in phase r each station n may store a different current phase number r_n. If in stage $k-1$ no station has played greedy or selfish, station n keeps selecting $w_{\rm h}$ in successive stages until greedy or selfish play is detected. It does so also upon detection of a downward trend in the observed play path, i.e., $x^{k-1} \prec x^{k-2}$. If selfish play has been detected in a stage, station n selects $w_{\rm s}$ with probability p^{r_n} and $w_{\rm h}$ with probability $1-p^{r_n}$ in successive stages until it detects either a downward trend in the observed play path, in which case it keeps selecting $w_{\rm h}$ again, or an upward trend, in which case phase $r+1$ starts (an event referred to as a *phase-up*) and the selfish/honest toggling continues. If station n detects greedy play in a stage, it selects $w_{\rm g}$ with probability p^{r_n} and $w_{\rm s}$ with probability $1-p^{r_n}$ in successive stages until it detects either a downward trend in the observed play path, in which case it retreats to selfish/honest toggling as described previously, or an upward trend, in which case the greedy/selfish toggling continues upon a phase-up.

Given the observed play path $(x^1, \ldots, x^{k-1})$, the next selection is determined by the pair $\boldsymbol{x}^{(k-1)} = (x^{k-2}, x^{k-1})$. A set of pairs $\boldsymbol{x}^{(k-1)}$ for which the selection follows in the same way defines a *CRISP state* prior to stage k. The five relevant CRISP states are defined by the following conditions (G, S, and H stand for greedy, selfish, and honest play):

- H: $x^{k-1} = 0$ or $x^{k-1} \prec x^{k-2}$,
- S/H: $0 \prec x^{k-1} = x^{k-2}$,
- S/H&PHASE-UP: $0 \prec x^{k-1}$ and $x^{k-2} \prec x^{k-1}$,
- G/S: $> M \prec x^{k-1} = x^{k-2}$, and
- G/S&PHASE-UP: $> M \prec x^{k-1}$ and $x^{k-2} \prec x^{k-1}$.

If σ_n = CRISP then, assuming that stage $k-1$ belongs to phase r and the probability distribution over W is encoded as ($Prob(w_n^k = w_{\rm g})$, $Prob(w_n^k = w_{\rm s})$, $Prob(w_n^k = w_{\rm h})$), we have

$$\sigma_n\left(x^{(k-1)}\right) = \begin{cases} (0,0,1), & \text{if } x^{(k-1)} \in \text{H}, \\ (0,p^{r_n},1-p^{r_n}), & \text{if } x^{(k-1)} \in \text{S/H}, \\ (0,p^{r_n+1},1-p^{r_n+1}), & \text{if } x^{(k-1)} \in \text{S/H} \\ & \quad \text{\& PHASE-UP}, \\ (p^{r_n},1-p^{r_n},0), & \text{if } x^{(k-1)} \in \text{G/S}, \\ (p^{r_n+1},1-p^{r_n+1},0), & \text{if } x^{(k-1)} \in \text{G/S} \\ & \quad \text{\& PHASE-UP}. \end{cases} \tag{13}$$

Each station starts off with an arbitrary $\boldsymbol{x}^{(0)}$; however, the x^k being public knowledge, from stage 3 on all the stations perceive a common CRISP state (and perform phase-ups in step). Fig. 1 depicts CRISP state transitions in stage k. The solid, dashed, and dotted arrows correspond respectively to the cases $x^k \prec x^{k-1}$, $x^k = x^{k-1}$, and $x^{k-1} \prec x^k$ and reflect the workings of CRISP without invaders. The transitions depicted as gray arrows marked I violate (13), hence can only occur in the initial stages and otherwise are solely due to the presence of invaders. Under all-CRISP, indefinite looping is possible within the states S/H and S/H&PHASE-UP, and within the states G/S and G/S&PHASE-UP, whereas H is an absorbing state corresponding to continued all-$w_{\rm h}$.

Fig. 2 shows an example scenario with $N = 3$ and $M = 1$; for clarity it is assumed that no station ever plays greedy. Stations 1 and 2 are standard, i.e., play CRISP according to (13); the phase-ups in stages 1 and 2 result from the initial states. Station 3 deviates by playing selfish in stage 6 despite $\boldsymbol{x}^{(5)} = (> M, 1, \ldots, M)$ (which would imply selecting $w_3^6 = w_{\rm h}$ under CRISP). Were it not for this deviation, all-$w_{\rm h}$ would settle

[3] A central result in game theory, the "folk theorem" [10], establishes the existence of such equilibria if the future and present stage payoffs are equally or almost equally weighted, which (11) does imply.

station	$\boldsymbol{x}^{(0)}$	phase 1	phase 2		phase 3				
	stage	1	2	3	4	5	6	7	8
1	$(>M, 0)$	w_{h}	w_{s}!	w_{h}	w_{s}!	w_{h}	w_{h}	w_{s}	w_{s}
2	$(1..M, >M)$	w_{s}!	w_{h}	w_{s}	w_{s}!	w_{h}	w_{h}	w_{h}	w_{s}
3	$(>M, >M)$	w_{h}	w_{h}	w_{s}	w_{h}!	w_{s}	w_{s}	w_{h}	w_{h}
	x^k	$1..M$	$1..M$	$>M$	$>M$	$1..M$	$1..M$	$1..M$	$>M$

Fig. 2. CRISP operation with three stations and $M = 1$ (shaded entries indicate configurations selected at random, ! indicates a phase-up).

in from stage 5 on. As it is, the two CRISP stations have entered the S/H state and continue to select their configurations at random, which will cause a phase-up when more than M select w_{s} (in our scenario, in nondepicted stage 9). Note why CRISP cannot rely on the outcome of the last stage only: for example, if $x^{k-1} \prec> M$ were to imply $w_n^k = w_{\mathrm{h}}$ then a single invader constantly selecting w_{s} would reap an unfairly large utility of $b_{\mathrm{s}}(N,1)$ (compare $s_{\mathrm{s}}(N,1)$ and $s_{\mathrm{h}}(N,0)$ in Table II).

Detection of a station selecting w_{g} brings about a painful punishment: all CRISP stations enter state G/S and require $x^k \prec \infty$ in order to stop toggling between w_{g} and w_{s}, which takes longer and longer to happen as r increases. Under light to moderate load, playing honest yields a distinctly nonzero bandwidth share, therefore $x^k = 0$ is constantly inferred (cf. Section III-B) and so all-w_{h} persists. The fact that CRISP only requires ternary granularity of x^k when no stations play greedy is of value given that the success probabilities may be observed (thus x^k may be inferred) inaccurately. Finally, as shown below, CRISP fulfills (12) for a wide class of p^r sequences and a range of M; this leaves room for trading the speed of payoff convergence to $b_{\mathrm{h}}(N,0)$ for punishment of invaders.

B. Pareto Efficiency and Subgame Perfection

We now look at sufficient conditions for CRISP to meet (12), assuming throughout that N remains constant during the game.

Proposition 2: Regardless of p^r, all-CRISP is fair and Pareto efficient.

Proof: We will show that for any initial play path, $u_n^k = (b_n^1 + \ldots + b_n^k)/k$ converges in probability to $b_{\mathrm{h}}(N,0)$. With the I transitions removed from Fig. 1, let random variable V represent the number of stages before reaching the absorbing state H. Since $b_n^k \equiv b_{\mathrm{h}}(N,0)$ for $k > V$, we can write

$$u_n^k = \begin{cases} b_{\mathrm{h}}(N,0) + (1/k)\sum_{i=1}^{V}\left(b_n^i - b_{\mathrm{h}}(N,0)\right), & \text{if } V \le k \\ (1/k)\sum_{i=1}^{k} b_n^i, & \text{if } V > k. \end{cases} \tag{14}$$

Recall that the stations may perceive different CRISP states in at most two initial stages. Therefore, it is convenient to write $V = V_3 + 2$, where V_3 is the number of stages before reaching H from the CRISP state in stage 3. Suppose that this state is S/H or S/H&PHASE-UP. Then $V_3 = V_3' + V_3''$, where V_3' counts the possible traversals of the lower left self-loop in Fig. 1 and the transitions between different CRISP states before reaching H, and V_3'' counts the traversals of the left self-loop, respectively. Clearly, V_3' is bounded (for $k \ge 3$ there can be two phase-ups at the maximum), whereas V_3'' either is geometrically distributed or is a sum of two independent geometrically distributed random variables (corresponding to play path segments with constant $x^k = 1..M$ and constant $x^k => M$). We can reason similarly if the CRISP state in stage 3 is G/S or G/S&PHASE-UP. Thus, $Prob[V \le k]$ tends to one as k increases. Since all the summands in (14) are bounded, the mean and variance of (14) tend to $b_{\mathrm{h}}(N,0)$ and zero, respectively. This proves our assertion, i.e., (12a) holds with $\sigma_n \equiv$ CRISP.

Let us now verify (12b). Clearly, it holds if σ prescribes playing honest indefinitely from some stage on, so assume that spells of honest play by σ are finite with probability one. That is, if l represents the length of a honest play spell then $Prob[l < \infty] = 1$, a "persistent invader" feature. Under the conditions stated below, σ yields a utility not exceeding $b_{\mathrm{s}}(N,N)$; this also proves that all-CRISP is a subgame perfect NE provided that $b_{\mathrm{s}}(N,N) < b_{\mathrm{h}}(N,0)$, which we have found true in existing IEEE 802.11 settings. □

Proposition 3: Let station n play a "persistent invader" σ, while $\sigma_m =$ CRISP for $m \ne n$. If $N - 1 > M$ and $\lim_{r\to\infty} p^r = 1$ then $U_n(\sigma_1, \ldots, \sigma_{n-1}, \sigma, \sigma_{n+1}, \ldots, \sigma_N|\pi) \le b_{\mathrm{s}}(N,N)$ for $\boldsymbol{\pi} \in \Pi$.

Proof: First consider the case when no station ever plays greedy. Assume (unrealistically) that station n knows *prior* to stage k the number x_{-n}^k of other stations playing selfish in stage k. Then there is no better play than to select $w_n^k = w_{\mathrm{s}}$ when $0 < x_{-n}^k \ne M$: not being able to cause a transition to another CRISP state, station n just maximizes its stage payoff according to part (ii) of Proposition 1. Let the resulting average stage payoff in phase r be $\overline{b}_n^r = A_n^r/Y^r$, where $A_n^r = b_n^{k_r} + \ldots + b_n^{k_{r+1}-1}$, $Y^r = k_{r+1} - k_r$ and k_r is the stage at the beginning of which the r^{th} phase-up occurs. Thus, it must be that $\boldsymbol{x}^{(k_r-1)} =$ S/H&PHASE$-$UP. In the case $x^{k_r-1} = 1..M$ phase r lasts one stage if $x^{k_r} => M$ (the lower left self-loop in Fig. 1 is traversed, i.e., another phase-up occurs), otherwise lasts a finite number of stages involving traversals of the left self-loop and possibly an ensuing spell of honest play by station n. Therefore, regardless of how station n replies to $x_{-n}^k = 0$, $\overline{b}_n^r = \sum_{x=M+1}^{N-1} \varPhi_{N-1}^r(x) b_{\mathrm{s}}(N, x+1) + (1-q^r)\overline{b}$, where $\varPhi_{N-1}^r(x)$ is the probability of x out of the $N-1$ CRISP stations selecting w_{s} in stage k_r, $q^r = \sum_{x=M+1}^{N-1} \varPhi_{N-1}^r(x)$, and $\overline{b}$ is a bounded value. As r increases, all the p^{r_m} tend to one; so does q^r and, since $N-1 > M$, also $\varPhi_{N-1}^r(N-1)$. (How station n responds to $x_{-n}^k = M$ becomes immaterial as r increases, since $\varPhi_{N-1}^r(M)$ becomes then arbitrarily small.) Hence, $\overline{b}_n^r$ converges in probability to $b_{\mathrm{s}}(N,N)$. If $x^{k_r-1} => M$ then with probability q^r (tending to one) phase r involves a nonzero number of traversals of the left self-loop. For large r, the mean of $\overline{b}_n^r$ conditioned on the number l of stages between a transition to state H and another phase-up can be expressed as

$$\begin{aligned} E\overline{b}_n^r &= \sum_{i=0}^{\infty} \frac{i\beta^r + l\overline{A}}{i+l}(q^r)^i(1-q^r) \\ &= \beta^r + l(\overline{A} - \beta^r)(1-q^r)\sum_{i=0}^{\infty}\frac{(q^r)^i}{i+l} \end{aligned} \tag{15}$$

where i is the number of traversals of the lower left self-loop, $\beta^r = (1/q^r)\sum_{x=M+1}^{N-1} \varPhi_{N-1}^r(x) b_s(N, x+1)$, and $\overline{A}$ is a bounded value. Since $Prob[l < \infty] = 1$, and recalling that

$\sum_{i=1}^{\infty} a^i/i = -\ln(1-a)$, we find that for large r the right-hand side of (15) asymptotes to β^r, hence tends to $b_{\rm s}(N,N)$ as r increases. Similarly, the conditional variance of $\overline{b}_n^r$ asymptotes to $l^2(\overline{A}-\beta^r)^2(1-q^r)\sum_{i=0}^{\infty}((q^r)^i/(i+l)^2)$, which tends to zero as r increases. Convergence in probability of $\overline{b}_n^r$ to $b_{\rm s}(N,N)$ has thus been established.

To complete the proof note that for large k and r, $u_n^k = (b_n^1 + \ldots + b_n^k)/k$ can be approximated by $U_n^r = (A_n^1 + \ldots + A_n^r)/(Y^1 + \ldots + Y^r) = (Y^1\overline{b}_n^1 + \ldots + Y^r\overline{b}_n^r)/Y_\Sigma^r$. Let r_0 be such that for all $r > r_0$, $Pr[|\overline{b}_n^r - b_s(N,N)| \le \varepsilon] \ge 1-\delta$. The number of occurrences $|\overline{b}_n^r - b_s(N,N)| > \varepsilon$ between r_0 and r (say in phases $r_1, r_2, \ldots$) is stochastically bounded[4] by a Bernoulli random variable $X_{r-r_0,\delta}$ with $r - r_0$ trials and success probability δ. Thus, with probability close to one, $|U_n^r - b_{\rm s}(N,N)|$ is

$$\frac{\left|Y_\Sigma^{r_0}\left[U_n^{r_0} - b_{\rm s}(N,N)\right] + \sum_{j=1}^{r} Y^j\left[\overline{b}_n^j - b_{\rm s}(N,N)\right]\right|}{Y_\Sigma^{r_0} + \sum_{j=1}^{r} Y^j} < \varepsilon + \frac{\overline{A}\left(Y_\Sigma^{r_0} + \sum_{i=1}^{X_{r-r_0,\delta}} Y^{r_i}\right)}{Y_\Sigma^{r_0} + \sum_{j=1}^{r} Y^j}$$

where $\overline{A}$ is a bounded value. It is visible that if ε and δ are sufficiently small and $r - r_0$ increases, the probability distribution of the right-hand side concentrates around zero. Thus, u_n^k converges in probability to $b_{\rm s}(N,N)$.

If greedy play is allowed then we distinguish two cases. If the invader plays greedy only finitely many times, the above reasoning can be repeated without modification. Otherwise for large r, phase r involves traversals of either the left or the right self-loop in Fig. 1 (in either case possibly followed by a spell of honest play). Thus, the mean of $\overline{b}_n^r$ is given, respectively, by (15) or an expression analogous to (15) with β^r replaced by 0 (the stage payoff in the presence of other stations playing greedy). Consequently, as k increases, the probability distribution of u_n^k concentrates below $b_{\rm s}(N,N)$. □

Strategies simpler than CRISP may not fulfill (12a) or (12b). Imagine an invader station n that never plays greedy, and a "deficient" CRISP, where $\boldsymbol{x}^{(k-1)} = (1..M, > M)$ is not followed by a phase-up. By always selecting $w_{\rm s}$, station n causes phase r to last forever, cycling within the states H, S/H, and S/H&PHASE-UP without increasing r. Then u_n^k is a weighted average of $b_{\rm s}(N,1)$ (when state H is passed) and $\sum_{x=0}^{N-1} \Phi_{N-1}^r(x) b_{\rm s}(N, x+1)$. This average may exceed $b_{\rm h}(N,0)$, so (12b) may be violated. As another example, suppose that no distinction is made between $x^k = 1..M$ and $x^k => M$, and only $x^k = 0$ and $x^k > 0$ are distinguished instead. To discourage an invader from always selecting $w_{\rm s}$, a phase-up should follow each pair $x^{k-2} > 0$, $x^{k-1} > 0$. Violation of (12a) is now possible since the succession of such pairs need not be finite with probability one. (Its *eventual* termination is guaranteed by the inverse Borel–Cantelli lemma [9] if the series $\sum_{r=1}^{\infty} \Phi_N^r(0) = \sum_{r=1}^{\infty}\prod_{n=1}^{N}(1-p^{r_n})$ diverges, e.g., if $p^r = 1 - r^{-\alpha/N}$ with $\alpha \in (0,1)$.)

[4]A random variable X is stochastically bounded by a random variable Y if $Prob[X \le k](Prob[Y \le k]$ for all k.

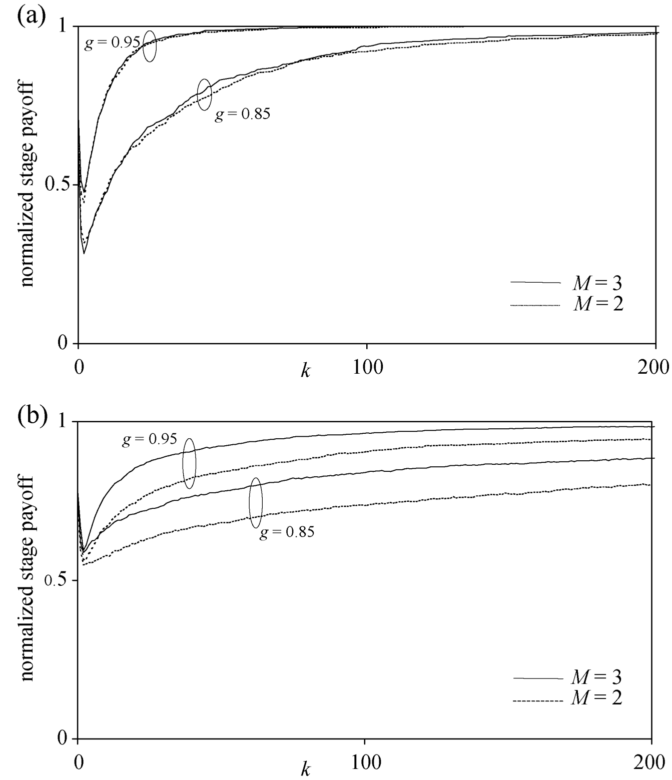

Fig. 3. All-CRISP convergence of stage payoffs to $b_{\rm h}(N,0)$: (a) small p^0; (b) large p^0 and no greedy play.

C. Performance

To illustrate the fulfillment of (12), the stage payoffs b_n^k were calculated from (5) for a 54-Mb/s IEEE 802.11a setting with basic access, 1500-byte DATA frames, $\tau_{\rm slot} \approx 61$, $\tau_{\rm DIFS} \approx 230$, $\tau_{\rm SIFS} \approx 108$, and $\tau_{\rm ACK} \approx 149$ (all durations expressed in byte transmission time units). Sample play paths of the repeated CSMA/CA game were generated via Monte Carlo simulation. Throughout the experiments, $N = 10$ was fixed. The sequence p^r was defined recursively as $p^{r+1} = 1 - g\cdot(1-p^r)$, where the parameter $g \in (0,1)$ controls the rate of growth of p^r as r increases. Four CRISP characteristics affect the performance: M, g, and the initial CRISP state $\boldsymbol{x}^{(0)}$ and probability p^0 at the start of a simulation run; these varied from experiment to experiment. Each CRISP station was initialized with a random p^0 chosen from (0, 0.2), (0, 0.5), or (0, 1), and with a random $\boldsymbol{x}^{(0)}$ chosen either from all the five feasible states or from H, S/H, and S/H&PHASE-UP only. Stage payoffs were normalized with respect to $b_{\rm h}(10,0) = 5.5\%$ and averaged over 1000 runs to produce satisfactory confidence intervals. In addition, the arithmetic average of the resulting averages was taken over all CRISP stations to produce a single representative CRISP payoff trajectory. By recording the average payoffs in successive stages one arrives at the plots in Fig. 3 and Fig. 4. Some reference normalized payoff levels are 12.1, 11.7, and 0.4, corresponding to the stage payoffs $b_{\rm g} = 69.9\%$, $b_{\rm s}(10,1) = 64.3\%$, and $b_{\rm s}(10,10) = 2.2\%$, respectively.

In Fig. 3(a) and (b), all the stations play CRISP, i.e., the I transitions do not occur (possibly except for the two initial stages). Eventually, each station detects a downward trend in the play

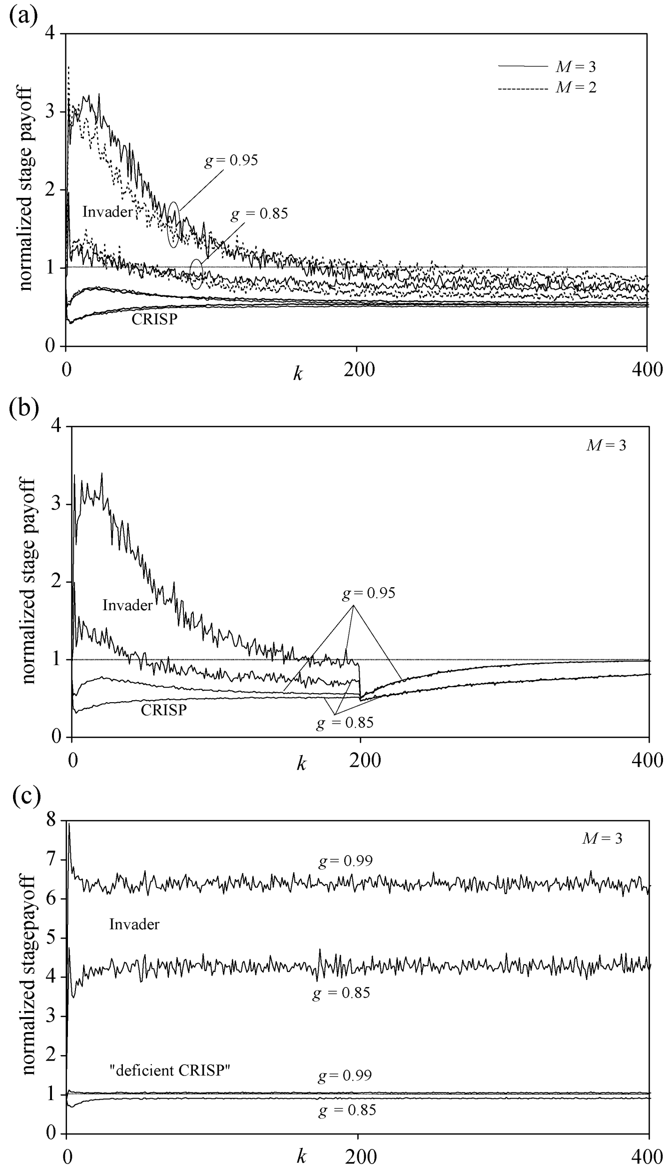

Fig. 4. CRISP in the presence of an invader: (a) punishment of "persistent invader"; (b) "learn and reform" scenario; (c) "persistent invader", versus "deficient" CRISP.

path and enters the absorbing state H. Thus, after a period of transient behavior, the configuration profile all-w_{h} sets in, with the stations' normalized stage payoffs approaching 1. This occurs the sooner, the larger are M and g ceteris paribus (note that g quantifies a station's reluctance to raise the probability of selecting w_{g} or w_{s} in successive phases). The relative significance of M and g depends on p^0 and $\boldsymbol{x}^{(0)}$. In Fig. 3(a), $\boldsymbol{x}^{(0)}$ is chosen from all feasible states and p^0 varies from 0 to 0.5. The $M = 2$ and $M = 3$ trajectories are barely distinguishable since in most cases the play path before all-w_{h} sets in only involves states G/S and G/S&PHASE-UP. A similar situation arises when $\boldsymbol{x}^{(0)}$ is chosen from H, S/H, and S/H&PHASE-UP only, which under all-CRISP implies that no station ever plays greedy. However, with p^0 chosen from (0, 1), a larger M visibly speeds up the convergence to $b_{\mathrm{h}}(N, 0)$, as it makes a downward trend in the play path more probable. E.g., with $g = 0.95$, all-w_{h} was observed in stage 50 in 53% and 84% of runs for $M = 2$ and $M = 3$, respectively, while with $g = 0.85$ the corresponding figures were 34% and 68%. For all-w_{h} to set in early, M and g should be configured jointly—taking a large M alone may not be the best option since the distinction of $x^k = 1, \ldots, M$ and $x^k \geq M$ might then be problematic, cf. Table II.

Not surprisingly, large M and especially g backfire if a "persistent invader" station is to be punished. In Fig. 4(a), all stations play CRISP except station n. In an attempt to conceive a "good" reply to CRISP when no station plays greedy, we let the invader predict $x_{-n}^k = |\{m \neq n | w_m^k = w_{\mathrm{s}}\}|$ prior to stage k and select $w_n^k = w_{\mathrm{s}}$ whenever $x_{-n}^k \neq M$ and $w_n^k = w_{\mathrm{h}}$ otherwise. If such play differs from what CRISP prescribes, the stage counts as a "selfish deviation" from CRISP. In the presence of greedy play, the invader predicts $y_{-n}^k = |\{m \neq n | w_m^k = w_{\mathrm{g}}\}|$ and selects $w_n^k = w_{\mathrm{g}}$ whenever $y_{-n}^k = 0$ and $w_n^k \neq w_{\mathrm{g}}$ otherwise. That is, it is careful not to prevent a downward trend in the play path, with the prospect of playing w_{g} in the next stage and landing a stage payoff of b_{g}. "Greedy deviations" from CRISP are likewise counted. We stress that such an invader strategy is unrealistic (unless station n is capable of instant accurate estimation of own and total success probability at the beginning of a stage). To take our idealization even further, we let station n optimize the number of "greedy deviations" and "selfish deviations" so that its payoffs stay above $b_{\mathrm{h}}(N, 0)$ as long as possible. In most experimented scenarios, optimum invader play consisted in no "greedy deviations" and infinitely many "selfish deviations," which confirms that invader stations may not be interested in greedy play, cf. the discussion in Section III-B. Nevertheless, one observes in Fig. 4(a) that the invader eventually fares worse than it would playing CRISP, its normalized stage payoffs slowly approaching 0.4. With large M and g the punishment for deviation from CRISP is less prompt and, within the first few dozens of stages, less severe. Clearly, the speed of convergence to all-w_{h} and punishment of invaders are in direct conflict and compromise CRISP parameters should be sought.

Fig. 4(b) illustrates the "learn and reform" environment created by CRISP. Station n plays the above invader strategy up to stage 200 and subsequently reverts to CRISP. As a result, both the invader's and the CRISP stations' normalized stage payoffs asymptote to 1 at a rate dependent on g. With $g = 0.95$, however, station n may have taken more than 200 stages to learn that deviation from CRISP is not beneficial. On the other hand, the subsequent convergence to $b_{\mathrm{h}}(N, 0)$ is much faster than for $g = 0.7$. (It would be even faster if station n had "reformed" before stage 200, e.g., having noticed the rapidly decreasing stage payoffs and inferred the presence of CRISP stations.)

As pointed out before, strategies less smart than CRISP may not resist smart enough invaders. For example, having failed against CRISP, the above idealized invader strategy adopted at station n turns out quite beneficial against a "deficient CRISP" in which no phase-ups occur [Fig. 4(c)]. When playing against $N - 1$ "deficient CRISP" stations with $M = 3$, station n observes the b_n^k hover in the region of four to seven times the fair payoff $b_{\mathrm{h}}(N, 0)$, the exact markup depending on M and g.

Finally, the postulate that CRISP be used as a standard strategy (i.e., standard stations are *committed* to CRISP) has a bearing on the CSMA/CA game: awareness of this is precisely what makes an invader "reform." Suppose to the contrary that standard stations are not committed to any strategy, but the invader is committed to the above described one, and that this is common knowledge. In Fig. 4(c), CRISP stations draw p^0 from (0, 0.2); this was found not to affect their payoffs significantly. Comparison of the plots in Fig. 4(b) and (c) reveals that for $M = 3$ and $g = 0.85$, "deficient CRISP" stage payoffs level off distinctly above those of regular CRISP stations (approximately 0.9 versus 0.5). Depending on g one may even expect "deficient CRISP" stage payoffs rise above $b_{\rm h}(N, 0)$ (cf. the $g = 0.99$ "deficient CRISP" trajectory, which levels off at about 1.05). That is, the presence of an invader may create a win-win situation. Without an incentive (in fact with a slight disincentive) to activate phase-ups, the standard stations' best reply is to let the invader enjoy its unfairly large bandwidth share indefinitely.[5]

D. Enforceability

Intuitively, any "persistent invader" will be asymptotically punished, and thus forced into reverting to CRISP, provided that enough stations play CRISP. This leads to the notion of *enforceability*.

Definition 7: Let σ^* fulfill (12) and σ be a "persistent invader" strategy. Strategy σ^* is *N, j-enforceable* if for any $\sigma_1, \ldots, \sigma_N$ with $|\{m|\sigma_m = \sigma^*\}| \geq j$, any $\boldsymbol{\pi} \in \Pi$, and any n such that $\sigma_n = \sigma$, $U_n(\sigma_1, \ldots, \sigma_N|\pi) < b_{\rm h}(N, 0)$.

That is, "persistent invaders" are punished if at least j out of N stations play σ^*. Note that Proposition 3 states conditions for $N, (N-1)$-enforceability of CRISP.

Proposition 4: Let

$$\chi(N,j,i) = \frac{\max_{1\le z\le i}\left[z\cdot b_{\rm s}(N,z+j) + (i-z)\cdot b_{\rm h}(N,z+j)\right]}{i}.$$

Assume that $\lim_{r\to\infty} p^r = 1$ and $j > \max\{M, \widetilde{j}(N)\}$, where $\widetilde{j}(N)$ is the largest integer j satisfying $\max_{1\le i\le N-j} \chi(N,j,i) \geq b_{\rm h}(N,0)$. There exists a $\varphi(N,j) < 1$ such that if $p^r \geq \varphi(N,j)$ for $r = 1, 2, \ldots$ then CRISP is N, j-enforceable.

Proof: Let j stations play CRISP, while i out of the other $N - j$ stations are "persistent invaders" never playing greedy. Let z^k of the latter play selfish in stage k, and suppose (unrealistically) that their total payoff is somehow distributed evenly among all the "persistent invaders," and that they can optimize z^k based on the predicted number x^k of stations to select $w_{\rm s}$ in stage k. An invader's utility is thus upper bounded by $\max_{1\le i\le N-j}\sum_{x=0}^{j} \Phi_j^r(x)\chi(N,x,i)$ (where $\Phi_j^r(x)$ was defined in the proof of Proposition 3). Consider first the case $N - j \leq M$. In each phase r, the number of CRISP state transitions is finite with probability one, since the probability of leaving the left self-loop in Fig. 1 is at least $\Phi_j^r(0) + \Phi_j^r(M+1)$ (the invaders are too few to cause indefinite looping with constant $x^k => M$ even if each of them selects $w_{\rm s}$ all the time and, since $j > M$, CRISP stations are sufficiently many to terminate looping with constant $x^k = 1..M$). Hence, the play undergoes infinitely many phases and so $\Phi_j^r(j)$ tends to one as r increases. In the case $N - j > M$, the invaders have the option of causing the CRISP stations to indefinitely self-loop with constant $x^k => M$ without a phase-up. If $p^{r_m} \geq \varphi$ then $\Phi_j^r(j) \geq (\varphi)^j$, therefore a φ close enough to one ensures that $\Phi_j^r(j)$ again becomes arbitrarily close to one. Thus, in either case, for φ above a threshold, the invaders each receive a utility not exceeding $\max_{1\le i\le N-j}\chi(N,j,i)$. Since $j > \widetilde{j}(N)$, this is less than $b_{\rm h}(N,0)$.

Greedy play by invaders does not affect N, j-enforceability. If any of them plays greedy infinitely many times then phases with greedy/selfish as well as with selfish/greedy toggling may occur. It is easy to see that in the former, indefinite looping in state G/S with exactly one station playing greedy occurs with probability zero; thus the probability distribution of stage payoffs concentrates around zero. □

[5]This may be viewed as an example of the so-called Stackelberg equilibrium [31], the invader assuming the role of the "leader."

TABLE III
CHARACTERISTICS RELATED TO CRISP ENFORCEABILITY

α	$j_\alpha(N)$ and $\widetilde{\varphi}(N, j_\alpha(N))$							
	$N=5$		$N=10$		$N=15$		$N=20$	
1	3	0.78	5	0.84	6	0.95	8	0.91
0.7	4	0.63	7	0.64	9	0.68	11	0.68

In light of the above argument, greedy play need not be considered a serious threat. The fact that $b_{\rm s}(N,N) < b_{\rm h}(N,0)$ in existing IEEE 802.11 settings makes Proposition 4 nontrivial, as it implies $\widetilde{j}(N) < N - 1$. To obtain a lower bound $\widetilde{\varphi}(N,j)$ on $\varphi(N,j)$ note that if $b_{\rm h}(N,x)$ and $b_{\rm s}(N,x)$ decrease in x (cf. Section III-B) then $\chi(N,x,i)$ does too. If $p^{r_m} \geq \varphi$ then by substituting $\Phi_j^r(x) = \binom{j}{x}(\varphi)^x(1-\varphi)^{j-x}$ into the inequality $\max_{1\le i\le N-j}\sum_{y=0}^{j}\Phi_j^r(x)\chi(N,x,i) < b_{\rm h}(N,0)$, one overestimates the left-hand side, so $\widetilde{\varphi}(N,j)$ is the minimum value of φ satisfying it. Denote $j_\alpha(N) = \min\{j|\widetilde{\varphi}(N,j) < \alpha\}$; in particular, $j_1(N) = \widetilde{j}(N) + 1$ is the minimum number of CRISP stations that can force the other stations into playing CRISP in an N-station WLAN, provided that the p^{r_m} are sufficiently large and that $M \leq \widetilde{j}(N)$. Table III shows $j_1(N)$ for several N; it follows that three CRISP stations can force up to two invaders (5,3-enforceability), five can force up to five (10,5-enforceability), six can force up to nine (15,6-enforceability), eight can force up to twelve (20,8-enforceability) etc. Roughly speaking, half the stations at the minimum should play CRISP to instill all-CRISP. The corresponding values of $\widetilde{\varphi}(N, j_1(N))$ are shown beside $j_1(N)$; they answer the question how to configure the p^r sequences at CRISP stations. Conversely, supposing that CRISP is configured with $p^r \geq 0.7$ (say), how many CRISP stations are necessary to instill all-CRISP given N? This amounts to finding $j_{0.7}(N)$; exemplary values appear in the bottom row.

VI. CONCLUSION

Given that a selfish station n can fix any w_n as an alternative to the standard $w_{\rm h}$, we have argued that $w_{\rm s}$ and $w_{\rm g}$ are reasonable choices. If a station can switch between $w_{\rm s}$ and $w_{\rm h}$, a CSMA/CA game arises that resembles a multiplayer Prisoners' Dilemma with a unique Pareto inefficient NE; if $w_{\rm g}$ is allowed, any NE is either Pareto inefficient or unfair and yields zero payoffs to most of the stations. As a defense against the backoff attack, a strategy called CRISP has been proposed and found to have a few desirable properties. It leaves the MAC protocol unaffected and only relies on stage-by-stage observations of success probabilities; the toggling between $w_{\rm g}$, $w_{\rm s}$, and

w_h can be performed by NIC driver software. It is not invoked under light or moderate load, i.e., when selfishness is not a real danger. Neither does it involve guessing the other stations' configurations or violation of station anonymity. Deviations from CRISP are discouraged on a disincentive basis in that a single invader station finds its bandwidth share asymptotically inferior to what it would receive playing CRISP. Finally, a certain minimum number of CRISP stations can force multiple invaders into playing CRISP. The inherent vulnerability to a malicious (denial-of-service) attack does not disqualify CRISP, given that such an attack can be launched anyway, e.g., using random jamming.

Compared with [7], CRISP attributes more introspection to the stations, leading to the conviction that only w_g, w_s, or w_h are likely to be configured. It also requires synchronization to stage boundaries. In return, no selective punishment for misbehavior is necessary (which in [7] permits to convert any configuration profile into a NE, but implies both station identification and violation of the MAC standard). Moreover, not all play paths of the dynamic game described in [7] end up at a Pareto efficient NE.

Subgame perfection of all-CRISP was only proved for the liminf-type utility (11). Thus, WLAN stations are assumed to put equal weight on near and distant future payoffs (i.e., there is no discounting of future payoffs). This sounds a little unrealistic in volatile WLAN environments and implies only asymptotic punishment of a "persistent" invader, whereas a "nonpersistent" one will not be punished at all. However, as Fig. 4(a) shows, punishment of a "persistent" invader typically follows within a reasonable time (on order of a few hundred stages). One can also see that CRISP is better than a trivial punishment strategy (e.g., jamming all DATA frames) for the standard stations still receive 50% to 70% of the fair bandwidth share while punishing an invader.

Two more drawbacks of CRISP require further attention. First, M must be configured relative to N, which is unknown to any station and may change over time. However, N can be inferred from $s_\mathrm{h}(N,0)$, cf. Table II, and dynamic agreement on M among CRISP stations can be envisaged. Second, the p^{r_n} have to be kept from permanently staying close to one; occasional reset to moderate values may be in order after a long enough period with no invaders detected.

ACKNOWLEDGMENT

The author would like to thank the anonymous referees for their valuable comments which helped to improve the presentation.

REFERENCES

[1] A. Akella, S. Seshan, R. Karp, S. Shenker, and C. Papadimitriou, "Selfish behavior and stability of the internet: a game-theoretic analysis of TCP," in *Proc. ACM SIGCOMM'02*, Pittsburgh, PA, Aug. 2002, pp. 117–130.

[2] E. Altman, R. El Azouzi, and T. Jimenez, "Slotted Aloha as a game with partial information," *Comput. Netw.*, vol. 45, pp. 701–713, 2004.

[3] R. Axelrod, *The Evolution of Cooperation*. New York: Basic Books, 1984.

[4] J. Bellardo and S. Savage, "802.11 denial-of-service attacks: real vulnerabilities and practical solutions," in *Proc. USENIX Security Symp.*, Washington DC, Aug. 2003, pp. 15–28.

[5] N. Ben Salem, L. Buttyan, J.-P. Hubaux, and M. Jakobsson, "A charging and rewarding scheme for packet forwarding in multi-hop cellular networks," in *Proc. ACM Symp. Mobile Ad Hoc Networking and Computing (MobiHoc 2003)*, Annapolis, MD, Jun. 2003, pp. 13–24.

[6] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, Mar. 2000.

[7] M. Cagalj, S. Ganeriwal, I. Aad, and J.-P. Hubaux, "On selfish behavior in CSMA/CA networks," in *Proc. IEEE INFOCOM 2005*, Miami, FL, Mar. 2005, pp. 2513–1514.

[8] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, "A BGP-based mechanism for lowest-cost routing," in *Proc. 21st Symp. Principles of Distributed Computing*, Monterey, CA, Jul. 2002, pp. 173–182.

[9] W. Feller, *An Introduction to Probability Theory and its Applications*. New York: Wiley, 1966.

[10] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.

[11] B. Grofman and J. Pool, "How to make cooperation the optimizing strategy in a two-person game," *J. Math. Sociology*, vol. 5, pp. 173–186, 1977.

[12] J.-P. Hubaux, L. Buttyan, and S. Capkun, "The quest for security in mobile ad hoc networks," in *Proc. ACM Symp. Mobile Ad Hoc Networking and Computing (MobiHoc 2001)*, Long Beach, CA, Oct. 2001.

[13] *IEEE Standard for Information Technology—LAN/MAN—Specific Requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, ISO/IEC 8802-11:1999.

[14] E. Kalai and E. Lehrer, "Rational learning leads to Nash equilibrium," *Econometrica*, vol. 61, pp. 1019–1045, 1993.

[15] V. Knoblauch, "Computable strategies for repeated Prisoner's Dilemma," *Games and Economic Behavior*, vol. 7, pp. 381–389, 1994.

[16] C. E. Koksal, H. Kassab, and H. Balakrishnan, "An analysis of short-term fairness in wireless media access protocols," in *Proc. ACM SIGMETRICS*, Santa Clara, CA, Jun. 2000, pp. 118–119.

[17] J. Konorski, "Multiple access in ad hoc wireless LANs with noncooperative stations," in *Proc. Networking 2002*, vol. LNCS 2345, pp. 1141–1146.

[18] ——, "Solvability of a Markovian model of an IEEE 802.11 LAN under a backoff attack," in *Proc. IEEE MASCOTS 2005*, Atlanta, GA, Sep. 2005, pp. 491–498.

[19] Y. A. Korilis, A. A. Lazar, and A. Orda, "Architecting noncooperative networks," *IEEE J. Sel. Areas Commun.*, vol. 13, no. 7, pp. 1241–1251, Jul. 1995.

[20] P. Kyasanur and N. H. Vaidya, "Detection and handling of MAC layer misbehavior in wireless networks," in *Proc. Int. Conf. Dependable Systems and Networks*, San Francisco, CA, Jun. 2003, pp. 173–182.

[21] A. B. MacKenzie and S. B. Wicker, "Game theory and the design of self-configuring, adaptive wireless networks," *IEEE Commun. Mag.*, vol. 39, no. 11, pp. 126–131, Nov. 2001.

[22] ——, "Selfish users in ALOHA: a game-theoretic approach," in *Proc. Vehicular Technology Conf. Fall 2001*, Atlantic City, NJ, Oct. 2001, pp. 1354–1357.

[23] P. Michiardi and R. Molva, "A game theoretical approach to evaluate cooperation enforcement mechanisms in mobile ad hoc networks," in *Proc. WiOpt'03*, Sophia-Antipolis, France, Mar. 2003.

[24] P. Obreiter, B. Koenig-Ries, and M. Klein, "Stimulating cooperative behavior of autonomous devices—an analysis of requirements and existing approaches," in *Proc. 2nd Int. Workshop Wireless Information Systems*, Angers, France, Apr. 2003.

[25] A. Orda, R. Rom Rand, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Trans. Networking*, vol. 1, no. 5, pp. 510–521, Oct. 1993.

[26] O. Queseth, "Cooperative and selfish behaviour in unlicensed spectrum using the CSMA/CA protocol," in *Proc. Nordic Radio Symp.*, Oulu, Finland, Aug. 2004.

[27] V. Ramaiyan, A. Kumar, and E. Altman, "Fixed point analysis of single cell IEEE 802.11e WLANs: uniqueness, multistability and throughput differentiation," in *Proc. ACM SIGMETRICS'05*, Banff, Canada, Jun. 2005, pp. 109–120.

[28] M. Raya, J.-P. Hubaux, and I. Aad, "DOMINO: a system to detect greedy behavior in IEEE 802.11 hotspots," in *Proc. MobiSys 2004*, Boston, MA, Jun. 2004, pp. 84–97.

[29] A. Rubinstein, "Finite automata play the repeated Prisoners' Dilemma," *J. Econ. Theory*, vol. 39, pp. 83–96, 1986.

[30] B. A. Sanders, "An incentive compatible flow control algorithm for rate allocation in computer networks," *IEEE Trans. Comput.*, vol. 37, no. 9, pp. 1067–1072, Sep. 1988.

[31] S. Shenker, "Making greed work in networks: a game-theoretic analysis of switch service disciplines," *IEEE/ACM Trans. Netw.*, vol. 3, no. 6, pp. 819–831, Dec. 1995.
[32] H. W. Shin and I. J. Wassell, "Application of game theory for distributed dynamic channel allocation," in *Proc. IEEE 55th Vehicular Technology Conf., Spring 2002*, Birmingham, AL, May 2002, pp. 404–408.
[33] S. Smale, "The Prisoners' Dilemma and dynamical systems associated to noncooperative games," *Econometrica*, vol. 48, pp. 1617–1634, 1980.
[34] V. Srinivasan, P. Nuggehalli, C. F. Chiasserini, and R. R. Rao, "Cooperation in wireless ad hoc networks," in *Proc. IEEE INFOCOM 2003*, San Francisco, CA, Mar./Apr. 2003, pp. 808–817.
[35] W. Stanford, "Symmetric paths and evolution to equilibrium in the discounted Prisoners' Dilemma," *Econ. Lett.*, vol. 31, pp. 139–143, 1989.
[36] A. Urpi, M. Bonuccelli, and S. Giordano, "Modelling cooperation in mobile ad hoc networks: a formal description of selfishness," in *Proc. WiOpt'03*, Sophia-Antipolis, France, Mar. 2003.
[37] 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*, New York, Jun. 2002, pp. 599–607.
[38] X. Yao, "Evolutionary stability in the n-person iterated Prisoners' Dilemma," *BioSystems*, vol. 39, pp. 189–197, 1996.
[39] E. Ziouva and T. Antonakopoulos, "CSMA/CA performance under high traffic conditions: throughput and delay analysis," *Comput. Commun.*, vol. 25, pp. 313–321, 2002.

Jerzy Konorski received his M.Sc. degree in electrical engineering from the Technical University of Gdansk, Poland, in 1976, and his Ph.D. degree in computer science from the Institute of Computer Science, Polish Academy of Sciences, Warsaw, in 1984.

He is currently with the Department of Information Systems, Gdansk University of Technology, where he teaches probability theory, operational research, and network theory, and conducts research in computer networking, performance evaluation, information systems, data transmission, and distributed systems. He has worked on a number of government and European projects, and authored over 30 papers published in international journals or conference records, and co-authored another 10. His current work focuses on the application of game theory to medium access control in wireless networks.