Discouraging Traffic Remapping Attacks in Local Ad Hoc Networks

Jerzy Konorski and Szymon Szott, *Member, IEEE.*

***Abstract*—Quality of Service (QoS) is usually provided in ad hoc networks using a class-based approach which, without dedicated security measures in place, paves the way to various abuses by selfish stations. Such actions include traffic remapping attacks (TRAs), which consist in claiming a higher traffic priority, i.e., false designation of the intrinsic traffic class so that it can be mapped onto a higher-priority class. In practice, TRAs can be executed in IEEE 802.11 ad hoc networks using the Enhanced Distributed Channel Access (EDCA) function. This attack is easy to perform yet hard to prevent. We propose a distributed discouragement scheme based on the threat of TRA detection and punishment. The scheme does not rely on station identities or a trusted third party, nor does it require tampering with the MAC protocol. We analyze an arising non-cooperative TRA game and find that under certain realistic assumptions it only incentivizes TRAs if they are harmless to other stations; otherwise the selfish stations are induced to learn that TRAs are counterproductive.**

***Index Terms*—Ad hoc networks, IEEE 802.11, EDCA, QoS, game theory, selfish behavior, traffic remapping attack.**

I. Introduction

A LOCAL ad hoc network can be modeled as a collection of stations contending for a single wireless channel in order to exchange data without a fixed transmission infrastructure, central supervision, admission control, or trusted network-wide security mechanisms. Each station hosts user applications that generate higher-layer traffic of various Quality of Service (QoS) requirements dependent on the traffic class. The task of the underlying network protocols is to map the higher-layer traffic class onto some network-defined traffic class so as to provide the required QoS at the lower layers. MAC-layer class-based QoS provisioning in ad hoc networks often uses a DiffServ-like approach [1] whereby a higher-layer (e.g., IP) packet carries information (e.g., a Distributed Services Code Point, DSCP) that determines the medium access rights of a data frame the packet is converted to. However, without dedicated security measures in place, this approach paves the way to various abuses by selfish stations that pursue a better QoS than prescribed by the higher-layer traffic class they generate. Such behavior can be referred to as a *traffic remapping attack* (TRA); it consists in claiming a higher medium access priority by false designation of the higher-layer traffic class so that it can be mapped onto a higher MAC-layer priority.

A perfect scenery for TRAs are IEEE 802.11 ad hoc networks using the Enhanced Distributed Channel Access (EDCA) function [2]. EDCA is a collision avoidance scheme which enables QoS differentiation among multiple traffic classes by improving over the well-studied mechanisms of its predecessor, the Distributed Coordination Function (DCF). In DCF, stations defer a transmission of a data frame until a specified-length DCF Interframe Space (DIFS) period of idle medium has been detected and subsequently back off for a random number of time slots, as defined by the contention window (CW), at the same time limiting the transmission opportunity (TxOP), i.e., the amount of transmitted data per single access instance. EDCA extends this by having each higher-layer traffic class mapped onto an access category (AC) characterized by its own CW and TxOP limits, as well as an Arbitration Interframe Space (AIFS) analogous to DIFS. By assigning different AIFS, CW, and TxOP, parameters to each AC it is easy to prioritize medium access, thus offering better QoS to higher-priority traffic. The defined ACs are, in order of decreasing priority, voice (VO), video (VI), best effort (BE), and background (BK).

EDCA enables a selfish station to, e.g., shorten the AIFS or CW, or expand the TxOP for the AC it is using, in order to step up the relative priority of its traffic. MAC parameter manipulation of this kind is feasible with contemporary wireless card drivers [3] and has been well-studied, cf. Section II. However, a TRA has advantages over MAC parameter manipulation: it brings the frame more resources along the whole end-to-end route instead of just in the attacker's single-hop vicinity (the impact of which is not considered here) and is much simpler to perform, as shown in the lower part of Fig. 1. A user application claims a higher AC for its best effort higher-layer traffic using packet mangling software (e.g., Linux *iptables*) to substitute the DSCP with a real-time DSCP in its IP packets. These packets, when converted to data frames at the MAC layer, will be remapped from the BE AC to the VO AC, and thus will achieve a much higher throughput than they are entitled to relative to other traffic. Clearly, an attacker station performing the TRA is likely to deteriorate the QoS provided to *honest* stations which designate their higher-layer traffic class correctly. A TRA can be easily executed using local packet mangling software, which does not require access to the wireless card driver. In particular, the AC designation of a data frame, i.e., the AC field in the MAC header, need not be tampered with. This implies that superficial detection schemes comparing the AC field with the higher-layer traffic

S. Szott is with the Department of Telecommunications, Faculty of Computer Science, Electronics and Telecommunications, AGH University of Science and Technology, Krakow, Poland, e-mail: {szott}@kt.agh.edu.pl.

J. Konorski is with the Faculty of Electronics, Telecommunications and Informatics, Gdansk University of Technology, Gdansk, Poland, e-mail: jekon@eti.pg.gda.pl.

Manuscript received May 23, 2013; revised October 1, 2013 and January 23, 2014.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

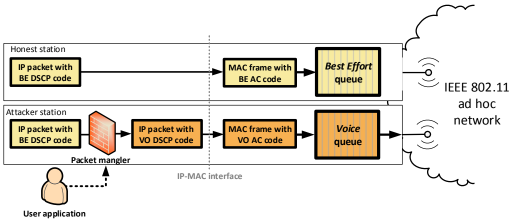
Honest station
IP packet with BE DSCP code
MAC frame with BE AC code
Best Effort queue
Attacker station
IP packet with BE DSCP code
Packet mangler
IP packet with VO DSCP code
MAC frame with VO AC code
Voice queue
IP-MAC interface
User application
IEEE 802.11 ad hoc network

Fig. 1. A traffic remapping attack using EDCA

designation, such as DSCP, are helpless against TRAs: unless the stations' individual throughput rates are closely monitored to detect unduly high throughput rates, the attacker can feel safe. TRAs can only be detected via costly traffic classification mechanisms (inspection of the frames' application data vis á vis their AC designation). In view of these difficulties, many proposed countermeasures against TRAs attempt to design incentives to properly designate generated higher-layer traffic (cf. Section II). These involve detection and punishment/reward based on a trusted third party (TTP) such as an access point (AP), and/or awareness of the number and identities of the contending stations.

In this paper, an extension of our previous conference papers [4] and [5], we consider a local ad hoc network with anonymous stations. Hence, an envisaged TRA detection via traffic classification (cf. Section VII.A) must not rely on station identities, and the punishment that may follow must be on a per frame, and not per source station, basis, e.g., using selective frame jamming [6]. Also, no TTP may be called upon. We take a game-theoretic approach, i.e., regard the selfish stations as (boundedly) rational players that strive to maximize some payoff in a non-cooperative game. We ask if the threat of detection can incentivize a selfish station to remain honest if the TRA could harm service levels obtained by other honest stations. At the same time, since it can improve the attacker's service level, traffic remapping ought to be allowed whenever it is harmless to other stations. Although we focus our analysis on EDCA, as practical means of executing a TRA, the derived conclusions can be applied to other distributed networks in which such attacks are a threat.

Each station is assumed to generate higher-layer traffic of one class only, hence can be identified by its 'type'—the AC its higher-layer traffic would be honestly mapped onto[1]. Thus one can speak of BE stations, VO stations, etc. In defining the payoff we assume that even for throughput-oriented BE stations the perception of received network services is binary rather than fine-grained: similarly as for VO or VI stations, there is a demanded throughput level above which the station is satisfied, and below which it is dissatisfied. Such an approach allows for a uniform treatment of all the ACs within the EDCA philosophy of service level provisioning. We show that if the payoff function combines the binary satisfaction from the received network services and the binary risk of (i.e., exposure to) detection and punishment, communicated through a simple broadcast scheme, then boundedly rational play leads to an operating point where all the stations are either honest or satisfied. That is, even if a TRA is performed by some stations, it does not diminish the satisfaction of other stations. Stations whose demanded service level is too high to be satisfied while remaining honest, and so are tempted to perform a TRA, will learn that it does not raise their payoffs in the long run. To summarize, we propose an incentive compatible soft security-type protection scheme against EDCA TRAs in an anonymous single-hop ad hoc network environment without a TTP.

The rest of the paper is organized as follows. In Section II we comment on existing related work. Section III is an overview of EDCA basics and sample effects of traffic remapping. In Section IV we formally define the TRA game in the context of EDCA and propose a simple payoff function based on the signaling of stations' dissatisfaction with received QoS or throughput. In Section V we present a simple, double-threshold way of playing the repeated TRA game conditioned on recently observed payoffs. We argue that such play reflects the stations' boundedly rational behavior and state the reachability of an 'all satisfied' operating point. Section VI illustrates via simulation that if the TRA game play in particular involves random choice of claimed AC then it (a) converges in probability to an 'all satisfied' operating point provided such a point is admitted by the demanded QoS and throughput levels, and (b) constitutes an equilibrium in that no deviation therefrom brings a distinct improvement of average payoffs. In Section VII we discuss the stations' capabilities of TRA detection, the observation of relevant information, and multi-hop considerations. Finally, Section VIII concludes the paper and outlines future work. The proofs of the assertions stated in the paper are provided in the Appendix.

II. Related Work

Selfish MAC-layer misbehavior, most notably backoff manipulation, has been shown to be both beneficial in terms of acquired bandwidth shares [7] and technologically feasible in some off-the-shelf equipment [3]. Various countermeasures have been studied for several years, mainly focusing on detection [8], [9] or prevention [10] of selfish misbehavior, or incentives for standard behavior [11]–[13]. Game-theoretic analyses (such as [14]) are particularly promising as they best model the autonomous and rational behavior of the wireless station, leading to self-regulatory, incentive compatible defense schemes; similar analyses also extend to more general wireless access settings [15]. So far, little attention has been given to multiple AC settings such as EDCA. Hu et al. [16] consider a game where an AP can admit or refuse a connection requested by a station, whereas a station can accept or reject the connection if admitted. This approach requires a trusted AP that knows the 'types' of all the stations, hence it does not address EDCA TRAs, which essentially consist in hiding a station's 'type' to claim a higher AC. Apparently, an EDCA TRA cannot be defended against using simple means. This is

[1] This clarifies the analysis, whereas considering stations sending traffic flows in multiple ACs does not provide much new insight into the behavior of the system as a whole. An honest multiple-AC station can be reflected in the analysis as multiple single-AC stations. For an attacker multiple-AC station it would be logical to send all traffic as the VO AC, hence become a single-AC station.

mostly because the 'type' is a station's private information, not necessarily revealed by the claimed AC. Cheung et al. [17] consider an infrastructure mode WLAN where stations are required to report their 'type' to the AP prior to data transfer, and to make payments to the AP for using a claimed AC. Within a mechanism design framework, the authors use a Vickrey-Clarke-Groves payment scheme to coax the stations into truthful 'type' reporting. This approach requires a trusted AP and a payment scheme implementation, which our solution does not need. In [18], Nguyen et al. restrict the considered ACs to one real-time and one bulk-data only. They note that using the real-time AC by all stations regardless of their 'types' is a Nash equilibrium of an AC selection game (an operating point at which no station can benefit by unilaterally changing the used AC), yet using the bulk-data AC brings higher throughput to stations sending QoS insensitive traffic. To resolve this dilemma, a modification of each AC's CW and TxOP limits is proposed so as to eliminate incentives for bulk data traffic to claim the real-time AC. However, it is assumed that a station can only decide on a (CW, TxOP) pair and not on individual parameters, which is not very practical. Our approach permits to retain the standard ACs and their defined parameters. Zhao et al. [19] consider using EDCA for prioritized transmission and an underlying priority selection game among selfish stations. To select an equilibrium strategy, each station must estimate the number of stations claiming each AC. By agreeing on the estimated numbers, the stations in fact engage in a cooperative game. Moreover, the proposed solution requires reliable station identities and an elaborate estimation scheme, which we dispense with. Li and Prabhakaran [20] studied a mixture of distributed and administrative decision making, where a station can relocate the priority of traffic incoming from a local application if it detects that its QoS requirements do not justify the claimed AC. It is thus assumed that MAC-layer devices are always honest and fully independent of the local applications (hence, users), which contradicts the usual conviction that a wireless station's equipment and users are united by a common goal. Our work does not distinguish between the goals of a station and its user. Nuggehalli et al. [21] propose a Hybrid Control Function (HCF) game framework for time slot assignment to two traffic classes; an efficient Nash equilibrium requires that a trusted supervisor (AP) only polls selected stations in the contention-free intervals based on their 'type' declarations (a similar approach is taken by Price et al. [22]). In contrast, our approach retains the fully distributed flavor of inter-station interactions and requires no supervisor. Ghazvini et al. [23] envisage a game in which, by setting appropriate TxOP values, stations maximize individual combinations of throughput and delay, weighted depending on a station's 'type'. The stations have distinguishable identities and can modify their MAC parameters, which our solution does not assume. Galluccio [6] takes an approach slightly similar to ours in that traffic flows performing a TRA face punishment, e.g., other stations might resort to selective frame jamming once they detect priority-related misbehavior. This prospect makes misbehaving stations act with restraint, so that the threat of punishment may never materialize. Unlike in our setting, stations' reliable identities, throughput estimation by priority, and knowledge of the number of stations claiming a given priority are required. Furthermore, we note that our solution concept for the proposed TRA game (the 'all satisfied' strategy profile) is in the spirit of satisfaction equilibrium, a concept introduced by Ross and Chaib-draa [24].

Based on the above analysis, we conclude that most existing studies on noncooperative MAC assume a single traffic class, making TRA meaningless. Those that deal with attacks analogous to TRA are functionally different from ours, e.g., by requiring a TTP, a payment scheme, static station identities, estimation of the number of competing stations, modifications of standard MAC, or extra signaling (cf. Table I). Thus, to the best of our knowledge, there exist no directly comparable solutions.

III. EDCA Operation And Performance

A user's requirements as to MAC performance of the ACs are strictly related to the higher-layer traffic class it generates. In the following, we consider each of the four EDCA ACs: VO, VI, BE, and BK, although in subsequent sections, to reduce the number of design variables, we only focus on the VO and BE ACs as they are expected to accommodate representative QoS- and throughput-sensitive traffic classes. Each AC has its characteristic MAC parameters as shown in Table II and discussed below.

VO and VI stations want their traffic to meet certain QoS requirements. Such requirements are not defined in the IEEE 802.11 standard, but can be found in the ITU-T Recommendation Y.1541 [25]. A mapping between EDCA ACs and ITU-T Classes of Service (CoS) can be found in [26]. This mapping is presented in Table II, which moreover shows the upper bounds for four QoS parameters: delay, jitter, packet loss ratio (PLR), and packet error ratio (PER). A station of a given 'type' is satisfied if the QoS parameters of its transmitted traffic are within the limits shown. Note that these QoS parameters are not directly observable at a sender station. We discuss QoS perception issues in Section VII.B.

BE stations are interested in achieving high throughput, which they can observe directly. Yet they too can set minimum requirements—the demanded throughput level that marks the boundary between satisfaction and dissatisfaction. The throughput demanded by a station is assumed to be externally imposed, e.g., by the supported user applications. Finally, BK stations set no QoS or throughput requirements.

To study the impact of EDCA TRAs, we used the ns-2.28 simulator [27] with an EDCA patch [28]. A single-hop IEEE 802.11 HR/DSSS (High Rate / Direct Sequence Spread Spectrum) network served five BE stations, each generating 2 Mb/s of constant bit-rate data traffic, and five VO stations, each generating a 320 kb/s high-quality constant bit-rate audio stream. Only the BE stations had incentives for TRAs, which they performed by claiming the VO AC. Table III presents simulation results for a varying number of attackers, stating the throughput achieved by the BE stations (normalized to their offered load) and the PLR for the VO stations; the latter turned out to be the critical QoS requirement.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

TABLE I
A QUALITATIVE COMPARISON OF THE STATE OF THE ART (MM – MODIFICATIONS OF STANDARD MAC , TTP – TRUSTED THIRD PARTY (AP), IDs – KNOWLEDGE OF NUMBER AND IDs OF PLAYERS, OH – SIGNALING OVERHEAD, BC – BACKWARD COMPATIBILITY WITH IEEE 802.11).

Scheme	Description	(Dis)incentive mechanism	MM	TTP	IDs	OH	BC
Li and Prabhakaran [20]	Control of user-declared priorities using cross-layer (IP-to-MAC) and inter-station information exchange	Priority re-allocation upon detection of TRA by user	Yes	No	No	Yes	No
Nuggehalli et al. [21], Price et al. [22]	HCF with benevolent AP, ALOHA-type contention	Only stations declaring low priority or perceived as truth-telling polled in contention-free period	No	Yes	Yes	No	No
Zhao et al. [19]	Cooperative diversification of CWs, TRA not addressed	Reduction of prescribed CW by low-priority station causes more collisions with high-priority ones	Yes	No	Yes	No	Yes
Cheung et al. [17]	ALOHA-type contention, stations adhere to AP-prescribed transmit probabilities	VCG payments to AP induce truthful 'type' reporting	Yes	Yes	Yes	Yes	No
Galluccio [6]	DCF backoff cheating for higher throughput, TRA not addressed	Frame jamming if throughput higher than expected.	Yes	No	Yes	No	Yes
Hu et al. [16]	Admission control by AP, assumes truthful 'type' reporting (no TRA attempt)	Acceptance of granted admission brings higher utility	No	Yes	Yes	Yes	No
Nguyen et al. [18]	Stations select from AP-prescribed (CW, TxOP) pairs	Honest selection of (CW, TxOP) pair improves relevant measure (throughput or delay)	Yes	Yes	Yes	No	No
Ghazvini et al. [23]	TxOP-setting game, stations value throughput and delay depending on 'type'	Deviation of TxOP from NE lowers utility	Yes	No	Yes	No	Yes
Our solution	Distributed search for maximum utility combining service level and exposure to TRA detection	Threat of frame jamming upon TRA detection	No	No	No	No	Yes

TABLE II
MAPPING BETWEEN EDCA ACs AND ITU-T Y.1541 CoS. TYPICAL EDCA PARAMETERS FOR IEEE 802.11 HR/DSSS PHY.

AC	CWmin/CWmax	AIFSN [slots]	TxOP limit	CoS	Delay [ms]	Jitter [ms]	PLR	PER
VO	7/15	2	3 ms	0	100	50	1×10^{-3}	1×10^{-3}
VI	15/31	2	1.5 ms	1	400	50	1×10^{-3}	1×10^{-3}
BE	31/1023	3	single frame	2, 3, 4	—	—	—	—
BK	31/1023	7	single frame	5	—	—	—	—

Three game-theoretic observations from Table III are not unlike those pertaining to DCF and its structural similarity to a Prisoners' Dilemma [13], [29]: (i) TRA is the BE stations' dominating strategy, hence the 'all attack' strategy profile is a Nash equilibrium, (ii) the Nash equilibrium is Pareto ineffective, since the BE stations achieve lower throughput and the VO stations experience higher PLR compared to the 'all honest' profile, and (iii) an honest BE station's throughput degrades when the number of attackers increases. However, one or two attackers do not disrupt the satisfaction of the VO stations, whose average PLR then remains within the upper bound given in Table II. This indicates that TRAs may be harmless under certain traffic conditions and required QoS parameters, hence need not lead to a Prisoners' Dilemma. Indeed, a simulation setting different from the above may yield different observations. Fig. 2 presents the BE station's relative throughput gain when performing a TRA by remapping to the VO AC in the presence of a variable number of VO stations. The incentive to perform the TRA vanishes as the number of VO stations increases; beyond a certain number, the throughput gain becomes negative. The reason is that the attacker experiences an increased frame collision rate when using the VO AC's smaller CW. Furthermore, Fig. 3 (based on the same simulation setting as Table III but with 10 VO stations to better illustrate the change in PLR) shows that a TRA performed by a BE station need not change the throughput of honest BE stations or PLR of honest VO stations, depending on the offered load, though it is clear that if a change does occur, it is never for the better. In the remainder of the paper we assume a scenario where there are incentives both to attack (a TRA improves an attacker's QoS parameters) and to discourage attacks (a station switching its status from attacker to honest does not worsen other stations' QoS parameters). This is stated more precisely in Section IV.B.

TABLE III
SIMULATION RESULTS FOR 10 STATIONS

No. of attackers	Throughput of a BE station as a fraction of offered load		Average PLR at a VO station
	Attacker	Honest	
0	—	0.38	0
1	1	0.223	0.6×10^{-3}
2	0.794	0.04	1×10^{-3}
3	0.486	0.015	22.7×10^{-3}
4	0.324	0.008	49.1×10^{-3}
5	0.225	—	85.9×10^{-3}

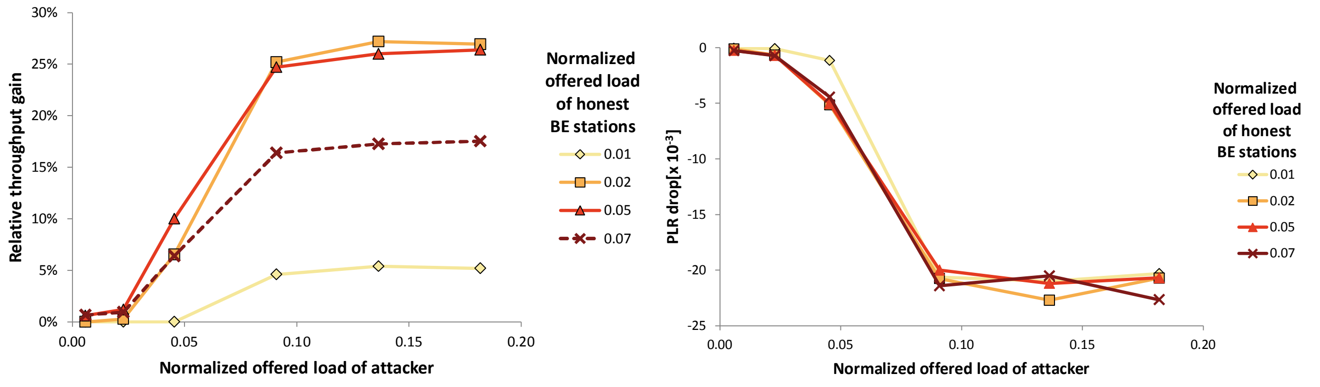

(a) (b)

Fig. 3. Relative throughput gain of an honest BE station (a) and PLR drop of a VO station (b) when another BE station switches from attacker to honest, as a function of the offered load normalized to the PHY data rate

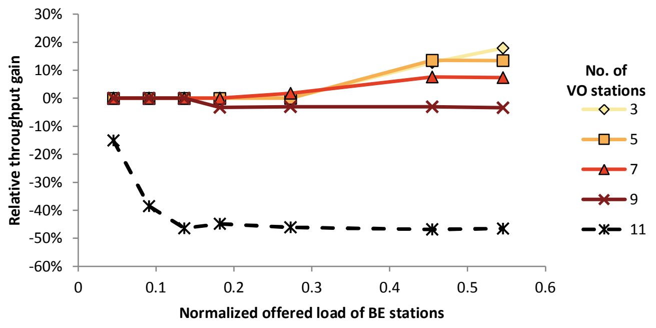

Fig. 2. Throughput gain of an attacker BE station in the presence of a varying number of VO stations, as a function of the BE stations' offered load normalized to the PHY data rate

IV. One-shot TRA Game

To formulate the problem of channel contention under IEEE 802.11 EDCA as a non-cooperative game we need to define the set of players, their feasible strategies, and the pay-off function. In the course of the game, the players choose their strategies autonomously and are not bound by any inter-player agreements that are not self-enforceable [29]. The set of feasible strategies is $C = \{\mathrm{VO}, \mathrm{VI}, \mathrm{BE}, \mathrm{BK}\}$, i.e., a strategy choice amounts to claiming an AC; by this we mean designation of generated higher-layer traffic so that it maps onto that AC. To reflect the induced medium access prioritization we number the ACs with their ITU-T CoS ordinals in reverse order, hence VO $>$ VI $>$ BE $>$ BK. Furthermore, a summary of the notation used in the paper is presented in Table IV.

A. Stations' Attributes

Consider N anonymous stations contending for a wireless channel under EDCA, each of which transmits traffic of one class only, i.e., has a well-defined 'type'. Denote by $\mathbf{r} = (r_1, r_2, \ldots, r_N)$, $\mathbf{t} = (t_1, t_2, \ldots, t_N)$, and $\mathbf{d} = (d_1, d_2, \ldots, d_N)$ the offered load, type, and demanded service level vectors, where for each station $i = 1, 2, \ldots, N$[2], r_i is the generated higher-layer traffic rate, $t_i \in C$ is the 'type', i.e., the AC that the higher-layer traffic class should be honestly mapped onto (all the local user applications are assumed to generate the same traffic class), and d_i is the demanded service level defined according to t_i. For a VO or VI station i, d_i represents the QoS parameter bound shown in Table II, whereas for a BE or a BK station, d_i is the demanded throughput (e.g., a BE station is obliged to serve a certain number of data connections with a guaranteed rate). It is reasonable to take $d_i = 0$ for a BK station and $0 < d_i < r_i$ for a BE station (since $d_i = 0$ would reduce it to a BK station, while $d_i = r_i$ would not leave room for data connections without a guaranteed rate). We assume the demanded service levels to be externally imposed. Thus we do not let station i 'play' with d_i in order to artificially perceive satisfaction or dissatisfaction with received network services. Both t_i and d_i are station i's private information; in view of the stations' anonymity, so is r_i.

A strategy profile is denoted by $\mathbf{c} = (c_1, c_2, \ldots, c_N)$ where $c_i \in C$ is station i's claimed AC. It is natural to assume that $c_i \geq t_i$, i.e., a selfish station presumes it will not benefit from claiming a worse service level than its 'type' would yield (this is justified by existing analyses of EDCA such as [7] and [16], as well as the simulations reported in Section VI (cf. also Definition 1 of Section IV.B). If $c_i = t_i$ then station i is honest, otherwise it is an attacker, i.e., performs an EDCA-based TRA. In our model, an attacker has a limited capability of reprogramming the wireless card. It can modify the MAC parameters of a chosen AC using the wireless card driver (though does not do this, as it expects to benefit from TRA); however, it cannot modify the MAC protocol itself. This allows us to focus on the problem of TRAs without combining it with other selfish attacks[3]. In particular, it implies that during a TRA there will always be a discrepancy between the AC field

[2] Station numbering is only used for notational convenience and does not imply any static identities.

[3] Such attacks include forcing a low priority AC frame to join a high priority AC queue. This attack is much harder to detect, since the DSCP remains genuine and AC handling is a station's private information. However, it is quite different from a TRA as no 'remapping' is performed, requires tampering with the MAC protocol, and only has a single-hop, rather than end-to-end scope. The threat of detection should in this case be strengthened by intelligent monitoring of the timing of sensed frames [30].

TABLE IV
SUMMARY OF NOTATION USED

c_i, $\mathbf{c}$, $\hat{\mathbf{c}}$	station i's claimed AC, strategy profile, and stationary strategy profile
C	set of feasible strategies (claimed ACs)
d_i, $\mathbf{d}$	station i's demanded service level and demanded service level vector
e_i	station i's exposure
F_i	station i's set of feasible strategies
k	stage of the repeated TRA game
l_i	station i's received service level
p_i, $\mathbf{p}$	station i's payoff and payoff vector
r_i, $\mathbf{r}$	station i's offered load (generated traffic rate) and offered load vector
s_i, $\mathbf{s}$	station i's satisfaction and satisfaction vector
t_i, $\mathbf{t}$	station i's traffic type and type vector
u_i^k, $\mathbf{u^k}$	station i's utility at stage k and utility vector at stage k
α_i	station i's learning coefficient
Γ	the one-shot TRA game
θ_{xi}, θ_{fi}	station i's explore and fallback thresholds

and user data contents of a frame, which is crucial for creating a threat of TRA detection as discussed below. Until such a detection is undertaken, c_i is station i's private information.

B. TRA Game and Stations' Payoffs

Given $\mathbf{r}$ and $\mathbf{t}$, the service level l_i received by station i is determined by $\mathbf{c}$, i.e., $l_i = l_i(\mathbf{c})$. Similarly as for d_i, l_i is defined according to t_i; for a VO or VI station it jointly represents the four QoS parameters in Table II, whereas for a BE station it represents received throughput. We take $l_i(\mathbf{c}) \geq d_i$ to mean that station i is satisfied with the received network service. With fixed $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$, a strategy profile $\mathbf{c}$ determines the satisfaction vector $\mathbf{s} = (s_1, s_2, \ldots, s_N)$, where

$$s_i = s_i(\mathbf{c}) = \begin{cases} 1 & \text{if } l_i(\mathbf{c}) \geq d_i \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

In accordance with the final remarks of Section III, in what follows we are interested in traffic and network conditions that incentivize both TRAs and TRA discouragement. We call such conditions a TRA game setting. Under fixed MAC parameters of standard ACs and temporal characteristics of typical higher-layer traffic, a TRA game setting may emerge depending on $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$.

Definition 1. $\mathbf{r}$, $\mathbf{t}$, *and* $\mathbf{d}$ *make a TRA game setting if*

(a) $s_i(\mathbf{c}) \geq s_i(\mathbf{t})$ *for some* $i = 1, 2, \ldots, N$ *and some* $\mathbf{c} > \mathbf{t}$ *with* $c_i > t_i$*, that is, at least one station may increase its satisfaction by performing a TRA, and*

(b) if $\mathbf{c}' < \mathbf{c}$ *with* $c_i' = c_i$ *then* $s_i(\mathbf{c}') \geq s_i(\mathbf{c})$ *for all* $i = 1, 2, \ldots, N$*, that is, if a subset of stations claim lower-priority ACs then the other stations' satisfaction does not decrease.*

While Fig. 2 indicates that part (a) of Definition 1 may not hold for some $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$, no simulations that we have carried out indicate violation of part (b). Yet one can suspect part (b) may not hold for some extremely impractical traffic and network conditions, which only accurate analysis could establish. Since existing EDCA analyses are approximate (Markov chain-based) and rarely lead to qualitative conclusions, we leave the issue open and for the rest of the paper assume the TRA game setting defined above.

The risk of (exposure to) detection of an EDCA-based TRA subtracts from an attacker station's satisfaction provided that the exposure is somehow communicated to the attacker and makes a credible threat of ensuing punishment. Moreover, only an honest station might be interested in ever triggering a detection and communication of the resulting exposure. For a network with anonymous stations, a simple signaling scheme can be proposed. A dissatisfied honest station (in practice, its dissatisfied recipient) broadcasts a DISSATISFACTION primitive appended to data or acknowledgment frames; this entitles any honest station that has sensed it to start a detection procedure (e.g., using traffic classification methods, cf. Section VII.A) and subsequently jam frames for which a TRA is detected. The detection may take a number of frames to discover a discrepancy between the AC field and user data contents, which gives enough time for attackers to retreat to honesty before a large number of their frames get jammed. In fact the very awareness of the exposure should prompt them to do so, assuming that they fear eventual punishment more than possible dissatisfaction when staying honest. Note that any abuse of the proposed scheme cannot adversely impact honest stations as these can afford to ignore DISSATISFACTION primitives. Hence, the use of traffic classification methods is a credible threat that may never have to materialize. Again considering $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$ as given, we formalize station i's exposure as follows:

$$e_i = e_i(\mathbf{c}) = \begin{cases} 1 & \text{if } c_i > t_i \wedge (\exists_{j \neq i} : c_j = t_j \wedge s_j(\mathbf{c}) = 0) \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Thus an attacker is only exposed in the presence of at least one station that is honest and dissatisfied at the same time. Finally, station i's payoff is defined as

$$p_i = p_i(\mathbf{c}) = s_i(\mathbf{c}) - e_i(\mathbf{c}), \tag{3}$$

implying that $p_i(\mathbf{c}) \in -1, 0, 1$. Given $\mathbf{r}$ and $\mathbf{d}$, a one-shot TRA game can be defined in the usual way as a triple specifying the set of players and their 'types', as well as the players' feasible strategy sets and the payoff function:

$$\Gamma = \langle \mathbf{t}; \mathbf{F}; p : \{1, 2, \ldots, N\} \times \mathbf{F} \to \{-1, 0, 1\} \rangle, \tag{4}$$

where $\mathbf{F} = F_1 \times F_2 \times \ldots \times F_N$ and $F_i = \{c \in C | c \geq t_i\}$. Note that VO stations can only choose the highest-priority VO AC, i.e., are always honest. Thus technically they are not players, although their dissatisfaction matters for other stations' exposure. Clearly, Γ is an incomplete information game, since the stations' 'types' and demanded service levels (thus also payoffs) are kept private. It is also an imperfect information game, since other stations' strategies are not observable.

Two kinds of strategy profiles in Γ are of particular interest: 'all honest', with $\mathbf{c} = \mathbf{t}$, and 'all satisfied', with $\mathbf{p}(\mathbf{c}) = \mathbf{1}$ (meaning $p_i(\mathbf{c}) = 1$ and implying $s_i(\mathbf{c}) = 1$ for all $i = 1, 2, \ldots, N$). Note that $\mathbf{p} = \mathbf{1}$ may or may not be admitted by $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$ (depending on whether the available bandwidth accommodates all the stations' demanded service levels). If $\mathbf{p} = \mathbf{1}$ is admitted then any strategy profile with $\mathbf{p} = \mathbf{1}$ is clearly a Nash equilibrium[4] of Γ and no profile with $\mathbf{p} < \mathbf{1}$ can be one. Thus Γ has multiple Nash equilibria, each possibly featuring a different number of attackers. If $\mathbf{p} = \mathbf{1}$ is not admitted then $\mathbf{t}$ ('all honest') is a Nash equilibrium. Indeed, if some VO stations are dissatisfied at $\mathbf{t}$, any dissatisfied BE station becomes exposed upon executing a TRA, so cannot increase its (currently zero) payoff; clearly, a satisfied BE station cannot either. If all VO stations are satisfied at $\mathbf{t}$ then some BE station must be dissatisfied. A TRA cannot make it satisfied without leaving some other (honest) station dissatisfied (otherwise $\mathbf{p} = \mathbf{1}$), hence the attacker cannot increase its payoff, as it becomes exposed. Depending on $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$, there may also be other Nash equilibria with $\mathbf{p} < \mathbf{1}$. E.g., in the setting of Table III, if $t_i = \text{BE}$ for $i = 1, \ldots, 5$, with $(d_1/r_1, \ldots, d_5/r_5) = (0.8, 0.7, 0.03, 0.03, 0.03)$, and $t_i = \text{VO}$ for $i = 6, \ldots, 10$, with $d_i = 0.1\%$ PLR, then $\mathbf{c} = (\text{VO}, \text{VO}, \text{BE}, \text{BE}, \text{BE}, \text{VO}, \ldots, \text{VO})$ is a Nash equilibrium with $\mathbf{p} = (0, 1, 1, 1, 1, 1, 1, \ldots, 1)$. In the next section, we introduce the repeated TRA game and argue that Nash equilibria with $\mathbf{p} < \mathbf{1}$ are of little interest.

V. Repeated TRA Game

Considering that BK stations do not have incentives to play the game Γ, and that VO and VI stations define payoffs in a qualitatively similar way, we now restrict our attention to a two-type game Γ with $C = \{\text{VO}, \text{BE}\}$. This restriction has little bearing upon further reasoning and permits to just distinguish honest and attacker stations without breaking up the latter depending on their 'types' and claimed ACs. Let there be N_{BE} BE stations and $N - N_{\text{BE}}$ VO stations (recall that each VO station i always chooses $c_i = t_i$).

While one expects that $\mathbf{r}$ and $\mathbf{d}$ remain unchanged in the longer term, each BE station i may find it beneficial at times to switch between $c_i = \text{VO}$ and $c_i = \text{BE}$, i.e., between playing honest or attacker. In our repeated TRA game model, the time axis is divided into *stages* such that in each stage $k = 1, 2, \ldots$, the game Γ is played. Values related to stage k will be superscripted k, e.g., $\mathbf{c}^k$ and $\mathbf{p}^k = \mathbf{p}(\mathbf{c}^k)$ are the strategy profile and stations' payoffs in stage k, respectively. In what follows, we associate the term *strategy* with a single stage, whereas the term *behavior* refers to the rule of next-stage strategy choice.

Before presenting the details of station behavior in the repeated TRA game, we present the proposed scheme in a concise form. After each stage, a station calculates an average of its recent payoffs. It then compares this value with two thresholds to determine its strategy for the next stage. A station may either keep its strategy unchanged, search for a better strategy, or retreat to honesty. This is explained in Section V.A, whereas in Section V.B we consider the reachability of stationary operating points.

A. Boundedly Rational Behavior by BE Stations

We recognize that being simple devices of limited computing power, BE stations exhibit only boundedly rational behavior [31], in particular, when choosing a next-stage strategy are unable to look far ahead and discount future payoffs, do not take into account the distant past, and may be slow to react to the recent past. Hence we assume that at the end of stage k, station i chooses c_i^{k+1} in a myopic way, with a view of a possibly high p_i^{k+1}. This choice is only based on the current moving average of recent payoffs, termed *utility*:

$$u_i^k = (1 - \alpha_i) \times u_i^{k-1} + \alpha_i \times p_i^k, \tag{5}$$

with $\alpha_i \in (0, 1)$ being station i's learning coefficient and $u_i^0 = 0$. Note that for $k > 0$, $u_i^k \in (-1, 1)$.

Bounded rationality and imperfect information dictate a BE station behavior that follows simple heuristic rules only based on observation of own payoffs, e.g., responsive learning [32]. In our model, we define two thresholds for a BE station i: the *explore* threshold $\theta_{xi} \in (0, 1)$ and the *fallback* threshold $\theta_{fi} \in (-1, 0)$. When u_i^k rises above the explore threshold, station i must have frequently enjoyed payoffs of 1 (satisfaction without exposure) in recent stages, hence is willing to keep its strategy unchanged—this is of course the less likely, the more throughput the station demands. In the case when u_i^k lies between the two thresholds, unconditional keeping of the present strategy is not recommended since there must have been spells of dissatisfaction and/or exposure recently; still, the utility being relatively high, those spells may not have prevailed and station i is free to explore F_i, i.e., arbitrate between VO and BE, in search of a better strategy. Finally, when u_i^k drops below the fallback threshold, payoffs of -1 (dissatisfaction combined with exposure) must have been frequent recently and the station is alerted to the imminent TRA detection and punishment; in this case the station retreats to honesty. This double-threshold behavior is illustrated in Fig. 4; a BE station i chooses c_i^{k+1} from the set of feasible next-stage choices:

$$F_i^{k+1} = \begin{cases} \{c_i^k\} & \text{if } u_i^k \geq \theta_{xi} \\ F_i = \{\text{VO}, \text{BE}\} & \text{if } \theta_{fi} \leq u_i^k < \theta_{xi} \\ \{t_i\} & \text{if } u_i^k < \theta_{fi}. \end{cases} \tag{6}$$

The only constraint on the arbitration between VO and BE in the middle line is that it must ensure that either strategy

[4] $\mathbf{c}$ is a (weak) Nash equilibrium of Γ [29] if $p_i(c_i, \mathbf{c}_{-i}) \geq p_i(c', \mathbf{c}_{-i})$ for all $i = 1, \ldots, N$ and $c' \in F_i$, where $\mathbf{c}_{-i} = (c_{-1}, \ldots, c_{i-1}, c_{i+1}, \ldots, c_N)$ is player i's opponents' profile.

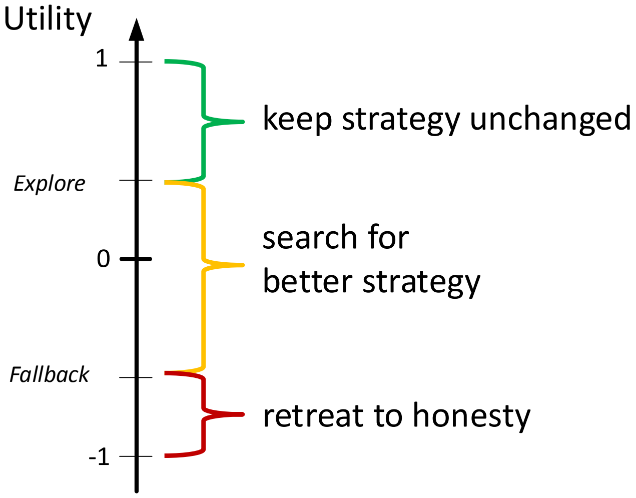

Fig. 4. BE station's next-stage strategy choice under (6)

is chosen infinitely often (i.e., infinitely many times in any infinite sequence of choices, though with no specific bounds on frequency). When considering various possible specifications of strategy choice from F_i^{k+1} one has to dismiss some obvious myopic rules, e.g., finding a best response c_i^{k+1} to the current opponents' profile $\mathbf{c}_{-i}^k = (c_i^k, ..., c_{i-1}^k, c_{i+1}^k, ..., c_N^k)$. Such behavior may in some cases lead to a Nash equilibrium of the game Γ, especially if the players are allowed to switch strategies one at a time, i.e., if $\mathbf{c}^k$ and $\mathbf{c}^{k+1}$ differ in at most one coordinate [33]. However, in an ad hoc network there is no way of enforcing any particular timing of strategy switching. Besides, station i is unable to observe c_j^k for $j \neq i$ (imperfect information), therefore, a best-response rule would be hard to define. In Section VI we show simulation results obtained for simple randomization between VO and BE.

B. Stationary Strategy Profiles and Their Reachability

Call a strategy profile $\hat{\mathbf{c}}$ *stationary* if it may persist forever when reached, i.e., under boundedly rational behavior (6) an infinite trajectory $\mathbf{c}^k = \mathbf{c}^{k+1} = \mathbf{c}^{k+2} = \ldots = \hat{\mathbf{c}}$ is feasible for some k.[5] With the above arbitration constraint, this implies that $F_i^{m+1} = \{c_i^m\}$ for all $m \geq k$ and $i = 1, 2, \ldots, N$. If the DISSATISFACTION scheme of Section IV.B is to discourage TRAs, the repeated TRA game must allow stationary strategy profiles that are desirable and, conversely, no undesirable strategy profile should be stationary. The latter postulate is addressed by the following assertion, stating that in the repeated TRA game, long-term TRAs are either precluded or harmless.

Assertion 1. *A stationary strategy profile is either 'all honest' or 'all satisfied' (if the latter is admitted by* $\mathbf{r}$, $\mathbf{t}$, *and* $\mathbf{d}$*), or both.*[6]

If an 'all satisfied' strategy profile, i.e., with $\mathbf{p}(\mathbf{c}) = \mathbf{1}$, is not admitted by $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$ then a persistent TRA should be discouraged in that it should not raise a delinquent BE station's utility (we address this postulate more precisely in the next section). However, if $\mathbf{p}(\mathbf{c}) = \mathbf{1}$ is admitted, it should be possible for boundedly rational stations "groping" for better strategies as prescribed by (6) to eventually happen upon such a profile $\mathbf{c}$. This conjecture is supported by the following assertion.

Assertion 2. *Suppose the repeated TRA game is in stage k and that* $\mathbf{r}$, $\mathbf{t}$, *and* $\mathbf{d}$ *admit a strategy profile* $\mathbf{c}^*$ *with* $\mathbf{p}(\mathbf{c}^*) = \mathbf{1}$. *Then there exists an 'all satisfied' strategy profile (possibly different from* $\mathbf{c}^*$*) that is reachable from* $\mathbf{c}^k$*, i.e., one* $\mathbf{c}^k$ *may transform into after a finite number of stages under (6).*

Note that if $\mathbf{c}^k$ is 'all satisfied', it will persist indefinitely if $u_i^k \geq \theta_{xi}$ for any BE station i. Then for large enough $m \geq k$, $\mathbf{u}^m \in \mathbf{1}^-$ ($\mathbf{1}^-$ symbolizes an arbitrarily narrow lower neighborhood of 1 in the N-dimensional Euclidean space). Such a course of play is desirable and turns out to be reachable, as stated below.

Assertion 3. *Suppose that* $\mathbf{r}$, $\mathbf{t}$, *and* $\mathbf{d}$ *admit an 'all satisfied' strategy profile. Then under (6),* $\mathbf{u} \in \mathbf{1}^-$ *is reachable from any stage of the repeated TRA game.*

The value of Assertion 3 is that it is robust to the threshold values θ_{fi} and θ_{xi} as long as the former is below and the latter is above 0. For any such setting, convergence to an 'all satisfied' profile, if one is admitted, is possible and even stochastically guaranteed for random arbitration in (6) (as demonstrated later in Section VI.A). This is convenient, since the thresholds are set autonomously at each station and are private information. Setting specific θ_{fi} and θ_{xi}, e.g., to optimize utility in the short run or speed up convergence to 'all satisfied', would require a station to peer into other stations' private information, namely their thresholds, offered loads, demanded service levels, and current strategies and payoffs. This seems impractical.

Referring to the final remarks of Section IV, Assertion 3 guarantees reachability from any initial setting, and asymptotic persistence, of an 'all satisfied' Nash equilibrium of Γ, if it exists. Since all such Nash equilibria are payoff equivalent ($\mathbf{p} = \mathbf{1}$), it makes little difference which one is asymptotically reached. If $\mathbf{p} = \mathbf{1}$ is not admitted, Nash equilibria other than $\mathbf{t}$ cannot be converged to (persist forever when reached), since by Assertion 1 they are not stationary profiles. The Nash equilibrium $\mathbf{t}$ is only converged to if in some stage k, $\mathbf{c}^k = \mathbf{t}$ and all the station utilities are below their respective fallback thresholds. The simulations in Section VI.B show this is rarely the case.

VI. Game Scenarios Under Random Arbitration

In this section we establish via simulation that under some additional specification, the play according to (6) converges to a stationary strategy profile and constitutes an equilibrium in that unilateral deviations from (6) bring no higher utility. To this end, we input the simulation results for the one-shot

[5] A stronger property can be postulated for a strategy profile: $\hat{\mathbf{c}}$ is *stable* if $\mathbf{c}^k = \hat{\mathbf{c}}$ implies $\mathbf{c}^m = \hat{\mathbf{c}}$ for all $m \geq k$, i.e., $\hat{\mathbf{c}}$ does persist forever when reached. However, from the proof of Assertion 1 one sees that under (5) and (6), no such $\hat{\mathbf{c}}$ exists in $\mathbf{F}$.

[6] The proofs of all assertions can be found in the Appendix.

two-type game Γ with $N = 10$ stations including $N_{\text{BE}} = 5$ BE stations and $N - N_{\text{BE}} = 5$ VO stations (Table III) as single-stage payoffs of the repeated TRA game, and next simulate the repeated TRA game play (6) to produce strategy profiles and station utilities in successive stages. For the simulations, the double-threshold behavior is specified as follows:

- the arbitration rule for the middle line of (6) consists in uniform (50%–50%) randomization between VO and BE; this is clearly in line with the arbitration constraint of choosing either strategy infinitely often, and
- the explore and fallback thresholds are related to their demanded throughput and offered load: $\theta_{xi} = d_i/r_i \in (0, 1)$, and $\theta_{fi} = d_i/r_i - 1 \in (-1, 0)$; coupling the two thresholds in this way reduces the number of design variables and ties the BE stations' play to their current satisfaction.

A. Convergence to Stationary Strategy Profiles

With uniform randomization between VO and BE in (6), the trajectory $(\mathbf{c}^0, \mathbf{u}^0), (\mathbf{c}^1, \mathbf{u}^1), \ldots, (\mathbf{c}^k, \mathbf{u}^k), \ldots$, where $\mathbf{u}^k = (u_1^k, u_2^k, \ldots, u_N^k)$, is governed by Markovian dynamics derived from (5) and (6):

$$\begin{cases} c_i^k = \mathrm{R}g_i(c_i^{k-1}, u_i^{k-1}) \\ u_i^k = (1 - \alpha_i) \times u_i^{k-1} + \alpha_i \times p_i(c^k) \end{cases} \tag{7}$$

for $i = 1, 2, \ldots, N$. In (7), $\mathrm{R}A$ is a randomly chosen element of a set A, and $g_i : F_i \times (-1, 1) \to 2^{F_i}$ is a set-valued function that, given a station's current strategy c_i^{k-1} and utility u_i^{k-1}, produces F_i^k according to (6) for a BE station i or $\{c_i^{k-1}\}$ for a VO station i. In this case, Assertion 3 can be strengthened by substituting stochastic convergence for reachability. Provided that an 'all satisfied' strategy profile is admitted, the only absorbing region of the above multi-dimensional Markovian dynamics is $\{(\mathbf{c}, \mathbf{u}) | \mathbf{p}(\mathbf{c}) = 1 \wedge \mathbf{u} \in \mathbf{1}^-\}$. Since reachability of this region is guaranteed by Assertion 3, so is stochastic convergence to it in the following sense: let $K(\mathbf{1}^-) = \min\{k | \mathbf{u}^k \in \mathbf{1}^-\}$; then $\Pr[K(\mathbf{1}^-) < \infty] = 1$, i.e., reaching $\mathbf{u} \in \mathbf{1}^-$ after a finite number of stages is almost certain [34].

The sample station utility and number of attackers' trajectories presented in Figs. 5 through 10 were obtained as averages from 20 simulation runs, each consisting of up to 2000 stages. The BE stations' learning coefficients appearing in (5) were chosen at random from the interval $[0.01, 0.2]$. To create a worst-case scenario from the viewpoint of reaching the 'all honest' or 'all satisfied' profile, all the BE stations were assumed to be attackers in the initial stage. In accordance with the remarks in Section V.B, no particular timing of strategy switching was enforced. (Other simulation experiments show that if the BE stations are only allowed to switch strategies one at a time, convergence to an 'all satisfied' profile, whenever admitted by $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$, is faster in some cases, albeit no qualitative difference is noticed. Hence, under randomization between VO and BE, the proposed TRA discouragement scheme is quite robust in producing a desirable outcome of the repeated TRA game.)

Given $\mathbf{r}$ and $\mathbf{t}$, as determined by the simulation setting of Table III, the key role in the distribution of station utilities rests with the demanded service level $\mathbf{d}$. For the VO stations we set $d_i = 1 \times 10^{-3}$ (in terms of PLR), so that a VO station remains satisfied in the presence of up to two attackers. Regarding the BE stations, we describe a few cases distinguished by the proportion of *aggressive* BE stations. A BE station i is called aggressive if d_i is set so that $l_i(\mathbf{t}) < d_i$, i.e., $s_i(\mathbf{t}) = 0$. That is, station i cannot be satisfied while honest even when all the other stations are also honest.[7] To increase its satisfaction, station i then has to persistently perform TRAs, which should eventually lower its utility unless no other station becomes dissatisfied in the process.

Fig. 5 presents a scenario with $d_i/r_i = 0.4$ for $i = 1, 2, \ldots, N_{\text{BE}}$. As seen from Table III, all the BE stations are aggressive and $\mathbf{p} = \mathbf{1}$ is not admitted. Half the BE stations on average perform TRA in each stage, receiving negative utilities. Though distinctly higher, the VO stations' utilities are also far below 1. This shows that when $\mathbf{p} = \mathbf{1}$ is not admitted, persistent TRAs by aggressive stations (i.e., those dissatisfied at $\mathbf{t}$, the 'all honest' profile) may preclude convergence to $\mathbf{t}$ and harm the honest stations. Yet these TRAs are even more harmful to the attackers, whose perceived risk of punishment increases. (Note that each aggressive BE station would receive zero utility if it was honest.)

In Fig. 6, $d_i/r_i = 0.22$ for $i = 1, 2, \ldots, N_{\text{BE}}$, which admits $\mathbf{p} = \mathbf{1}$ (no stations are aggressive). $\mathbf{u} \in \mathbf{1}^-$ is reached within the first 150 stages, whereupon the 'all satisfied' profile prevails. Yet the number of attackers stabilizes around one, since $\mathbf{p} = \mathbf{1}$ can accommodate up to one attacker (i.e., single-attacker TRAs are harmless). The scenario in Fig. 7 still features no aggressive stations, but the BE stations' demanded service level has risen to 0.23. This subtle increase has resulted in slowing down the convergence dramatically; the reason is that $\mathbf{p} = \mathbf{1}$ now cannot accommodate any attackers (recall from Table III that the critical demanded service level for the honest BE stations is 0.223), which takes the BE stations much longer to learn.

In Figs. 8 through 10, BE stations 1 and 2 are aggressive. We wish to emphasize the role of the aggressive stations' demanded service level, therefore we set $d_i/r_i = 0.022$ for $i = 3, 4, 5$ so that the nonaggressive BE stations can be satisfied with up to two attackers (recall from Table III that the critical demanded service level is 0.04). In Fig. 8, $d_i/r_i = 0.5$ for $i = 1, 2$. As seen from Table III, $\mathbf{p} = \mathbf{1}$ is admitted (the critical demanded service level for the aggressive BE stations being 0.794), and so the utilities of the nonaggressive BE stations rapidly converge to 1, as do the utilities of the VO stations. Not so rapid is the convergence for the two aggressive stations, whose distinctly higher demands postpone receiving payoffs of 1 and so prevent their utilities from quickly rising above the explore threshold. Since $\mathbf{p} = \mathbf{1}$ requires both of them to attack, the number of attackers approaches two. When d_i/r_i rises to 0.7 for $i = 1, 2$ (Fig. 9), $\mathbf{p} = \mathbf{1}$ is still admitted, yet convergence to $\mathbf{u} \in \mathbf{1}^-$ is now dramatically slower, due

[7] Being unable to observe the current strategy profile, a BE station is not aware of its aggressive status; this notion only serves to better understand the obtained utility trajectories.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

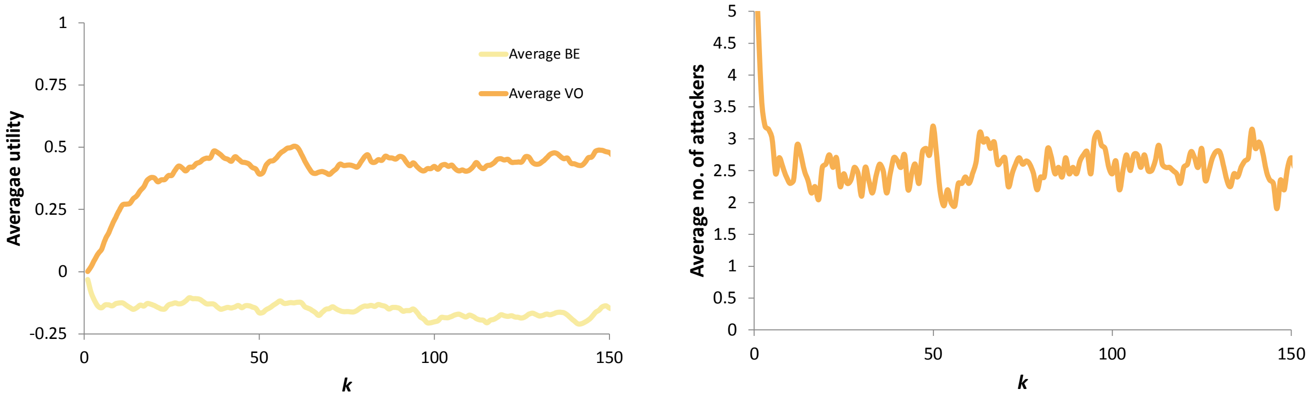

Fig. 5. Stations' utilities (*left*) and the number of attackers (*right*); all BE stations are aggressive, $\mathbf{p} = \mathbf{1}$ is not admitted.

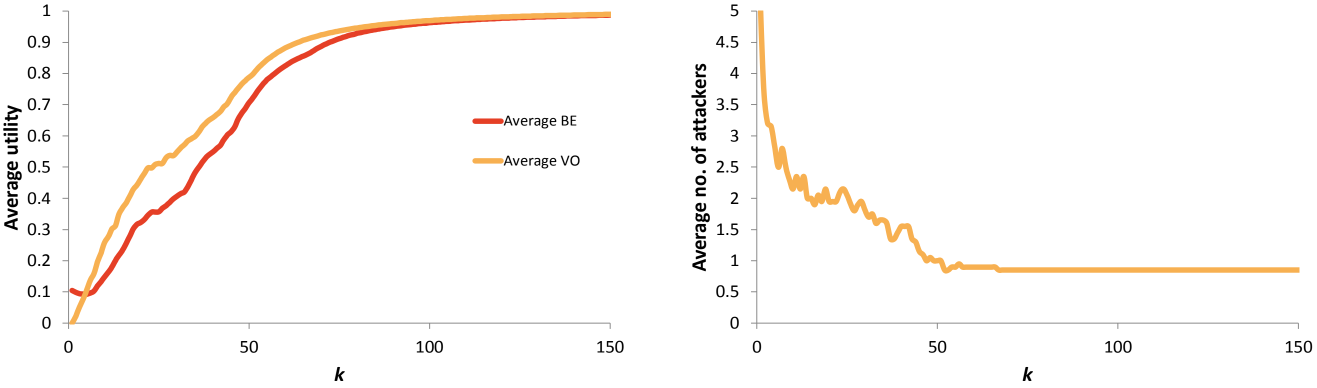

Fig. 6. Stations' utilities (*left*) and the number of attackers (*right*); no aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, rapid convergence.

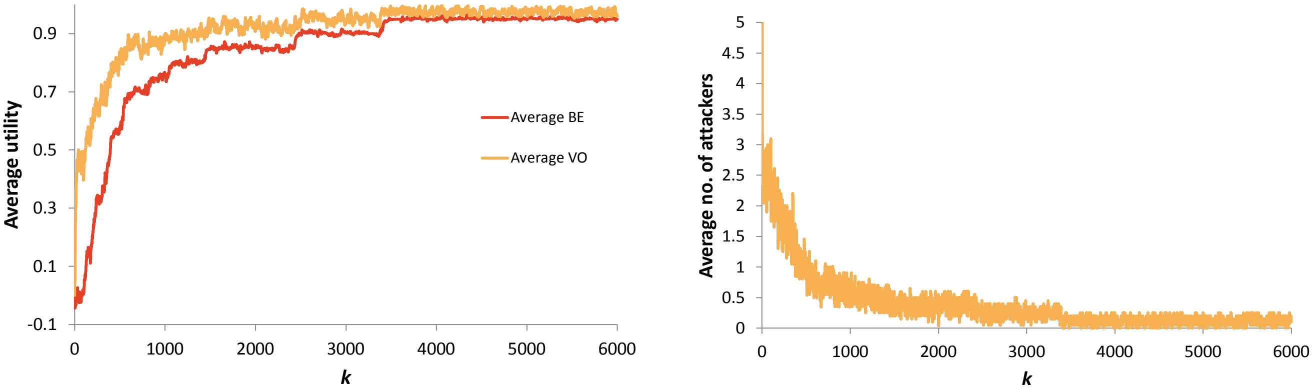

Fig. 7. Stations' utilities (*left*) and the number of attackers (*right*); no aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, slow convergence

to the even higher explore threshold at the two aggressive stations.

Finally, in Fig. 10, $d_i/r_i = 0.9$ for $i = 1, 2$, which is beyond the critical demanded service level for $\mathbf{p} = \mathbf{1}$ to be admitted. Stations 1 and 2 now achieve near-zero asymptotic utilities, distinctly lower than the other BE and VO stations. Note that unlike in Fig. 5, only the aggressive stations suffer a utility reduction. The number of attackers remains slightly below one, revealing that each of the aggressive stations performs TRAs persistently and in vain, but sometimes resorts to being honest, due to the high fallback threshold ($d_i/r_i - 1 = -0.1$). Meanwhile, all the nonaggressive BE stations quickly learn to play honest and become satisfied; so do the VO stations, albeit at a much slower pace. When stations 1 and 2 get discouraged and retreat to honesty, they will be dissatisfied of course, hence no station with a high demand will be satisfied in the presence

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

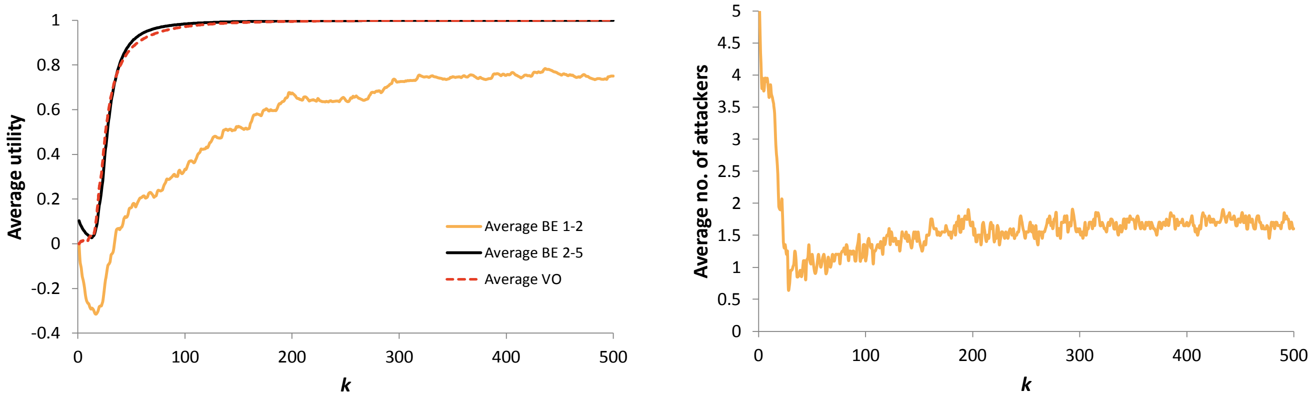

Fig. 8. Stations' utilities (*left*) and the number of attackers (*right*); two aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, rapid convergence.

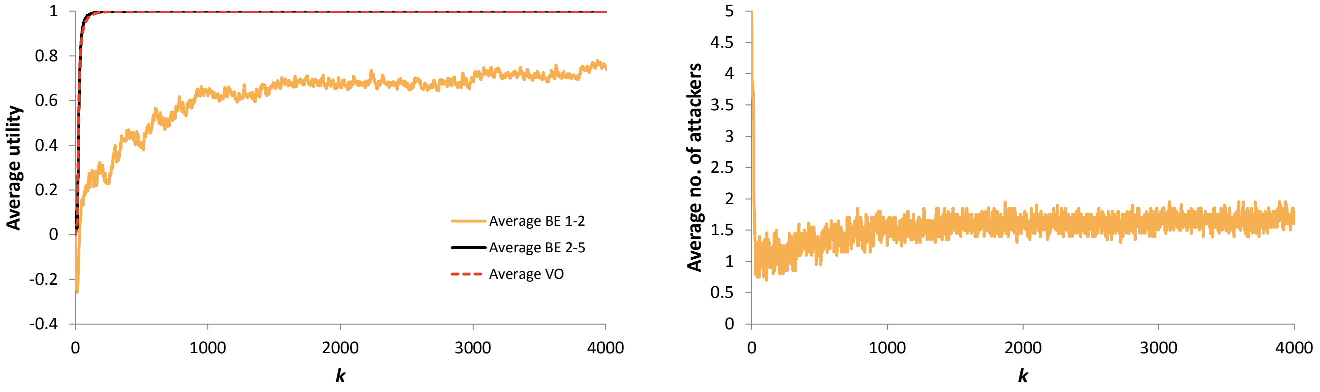

Fig. 9. Stations' utilities (*left*) and the number of attackers (*right*); two aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, slow convergence

of a dissatisfied station with a lower demand. This is in line with the max-min fairness philosophy [35].

In addition to the above studies, we have performed simulations for networks of various size and proportion of VO and BE stations to find that in general:

- stations are able to learn if a TRA is necessary to reach satisfaction and harmless enough for other stations to allow it,
- the proposed scheme leads to desirable strategy profiles (in which stations are 'all honest' or at least 'all satisfied') in any network configuration,
- an 'all satisfied' profile, if admitted, is always reached after a time that increases with the number of stations, the number of aggressive stations, and the average demanded throughput of the BE stations,
- in most cases, the convergence time is proportional to the learning coefficient α,
- the average number of attackers either converges to the maximum admissible number of attackers or approaches around half the total number of aggressive stations,
- the scheme tends to protect the payoffs of honest stations in the presence of aggressive stations in that convergence to $\mathbf{u} \in \mathbf{1}^-$ is faster for honest stations, and
- the convergence time could be decreased if additional factors, such as the monotonicity of recent u_i^k values, were included in the next-stage choice (6).

B. Behavior Deviation

We have shown that desirable strategy profiles can be reached under the double-threshold behavior specified by (5) and (6). If, in addition, no station deviating therefrom at some moment can expect an increased asymptotic utility regardless of the current history of play, the behavior is called subgame perfect [29]. For a rigorous assessment of subgame perfection in the repeated TRA game one should show that any boundedly rational behavior constituting a best response to (6) does not promise a higher utility, which is not easy; a frequent approach based on payoff discounting is not suitable here, as it does not fit in the bounded rationality paradigm. We offer partial insight by first noting that if $\mathbf{p} = \mathbf{1}$ is admitted then Assertion 3 in principle states subgame perfection, since regardless of the history of play, a deviator from (6) cannot expect a higher utility. The following can moreover be proved.

Assertion 4. *Suppose that (a) $N_{BE} < N$, (b) $\mathbf{p} = \mathbf{1}$ is not admitted, (c) BE station 1 is aggressive, and (d) some honest station different from station 1 cannot be satisfied in the presence of more than one attacker. Then under (6), $p_i(\mathbf{c}) \leq 0$*

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

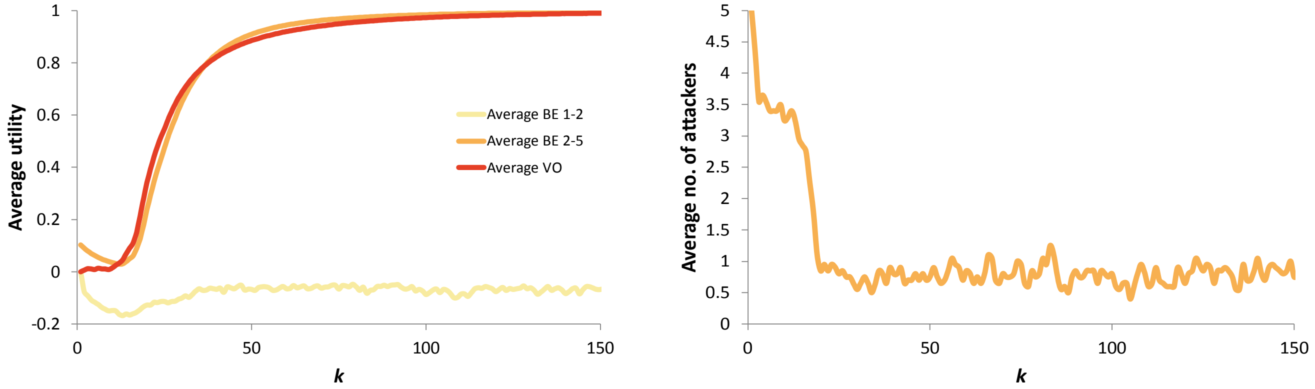

Fig. 10. Stations' utilities (*left*) and the number of attackers (*right*); two aggressive stations, $\mathbf{p} = \mathbf{1}$ is not admitted.

for all $\mathbf{c} \in \mathbf{F}$*, i.e., an aggressive station cannot receive a positive payoff.*

Thus if assumptions (a)-(d) above are fulfilled, an aggressive station receives nonpositive payoffs in successive stages and even the most intelligent deviation from (6) cannot yield a utility exceeding that of a persistently honest station, i.e., 0. Under (6), an aggressive station's asymptotic utility is not less than $\theta_{fi} - \alpha_i \times (1 - |\theta_{fi}|)$ (obtained by putting $u_i^{k-1} = \theta_{fi}$ and $p_i^k = -1$ in (5)). In conclusion, if $|\theta_{fi}| << 1$ and $\alpha_i << 1$ then there is little room for improvement of asymptotic utility by deviating from (6) regardless of the history of play.

Assumption (d) above is quite realistic, e.g., in the setting of Table III, it is enough that the BE stations demand a throughput larger than 4% of the offered load and/or the VO stations demand a PLR less than 0.1%. If Assertion 4 does not apply, one can assess subgame perfection of (6) via simulation similarly as in [13], by introducing an *ideal deviator* station capable of peering into other stations' private information and predicting their next-stage strategies. Although such behavior is impossible to implement in practice, it serves as a convenient reference. Let BE station 1 be the ideal deviator and assume it chooses its next-stage strategy as follows:

$$c_1^{k+1} = \begin{cases} \text{VO}, & \text{if } p_1(\text{VO}, \mathbf{c}_{-1}^{k+1}) > p_1(\text{BE}, \mathbf{c}_{-1}^{k+1}) \\ \text{BE}, & \text{otherwise.} \end{cases} \quad (8)$$

That is, station 1 chooses the strategy that yields a higher payoff against the strategy profile to be played in the next stage. Prior to the beginning of the next stage, station 1 must know the strategies chosen by the other players (which violates causality), their types, offered load and demanded service levels (which are all private information), as well as the service levels they are to receive (which would require instant analytical calculations even if the other pieces of information were known).

Figs. 11 through 14 present simulation results for the scenarios reported previously in Figs. 5, 8, 9, and 10, respectively. Station 1, unable to reach a utility of 1 in all those scenarios, is now the ideal deviator exhibiting the behavior (8). In all presented cases, the difference in utility between station 1 and the other BE stations is asymptotically marginal. In Fig. 11, where all the BE stations are aggressive, station 1 is aware that $\mathbf{p} = \mathbf{1}$ is not admitted. It is therefore able to maintain a utility of 0 (the highest possible utility for aggressive BE stations in this scenario). The other BE stations, only relying on locally available information, attempt TRAs, which results in a utility slightly below 0. Hence, there is little room for utility improvement that the behavior (8) promises and it certainly does not increase station 1's satisfaction.

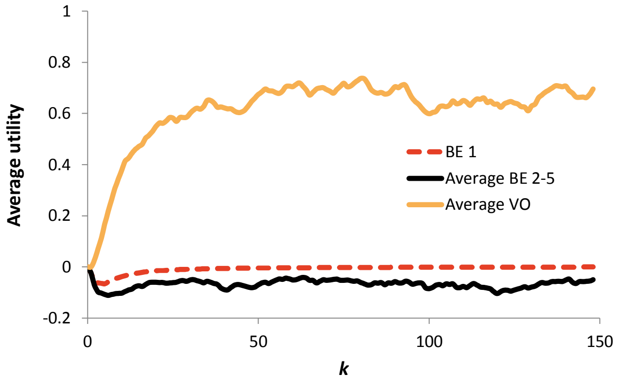

Fig. 11. Stations' utilities; all BE stations are aggressive, $\mathbf{p} = \mathbf{1}$ is not admitted, station 1 is the ideal deviator.

In the next figures, stations 1 and 2 are aggressive and behave according to (8) and (6), respectively. As long as the aggressive stations' demanded service level (whether low, as in Fig. 12, or high, as in Fig. 13) admits $\mathbf{p} = \mathbf{1}$, both these stations' utilities ultimately converge to 1. A difference in their utility trajectories is only present in the initial stages, when all the BE stations are in their learning phase. An additional effect caused by the presence of the ideal deviator is a more rapid convergence in comparison with Figs. 8 and 9.

When the aggressive stations' demanded service level rises beyond the critical level for $\mathbf{p} = \mathbf{1}$ to be admitted (Fig. 14), both stations ultimately reach a zero utility. Again, the only difference in their utility trajectories can be observed in the initial stages. However, after these initial stages and in

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

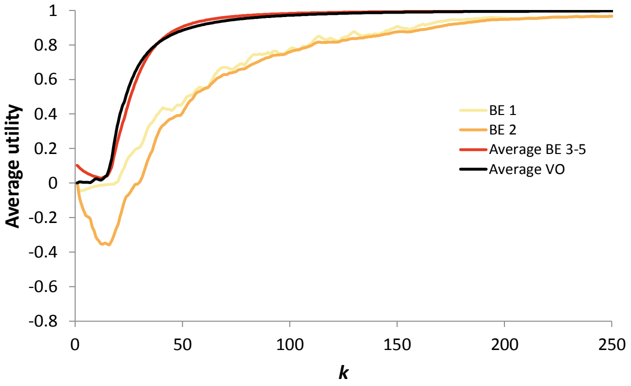

Fig. 12. Stations' utilities; two aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, rapid convergence, station 1 is the ideal deviator

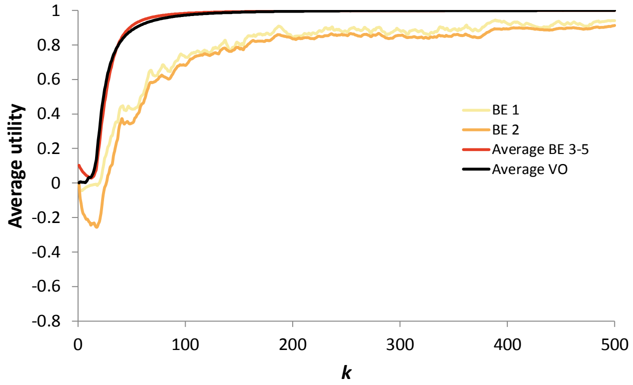

Fig. 13. Stations' utilities; two aggressive stations, $\mathbf{p} = \mathbf{1}$ is admitted, slow convergence, station 1 is the ideal deviator.

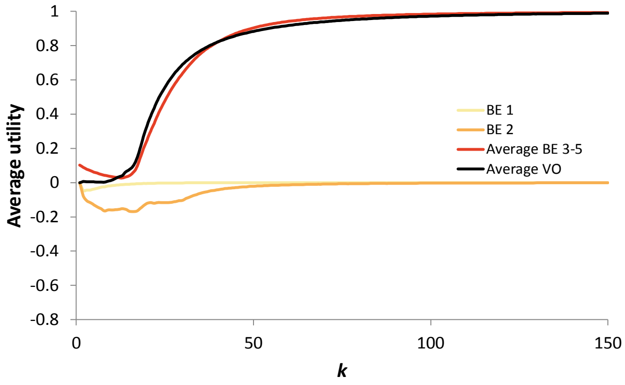

Fig. 14. Stations' utilities; two aggressive stations, $\mathbf{p} = \mathbf{1}$ is not admitted, station 1 is the ideal deviator.

comparison with Fig. 10, the utility of both aggressive stations is more stable. This is because (8) dictates that the ideal deviator always be honest, as it cannot gain from TRAs; on the other hand, station 2 keeps "groping" for better strategies, in successive stages being either honest and dissatisfied or an exposed attacker, and in both cases receiving zero payoffs.

Based on Assertion 4 and the above sample simulations we conclude that there do not exist deviations from the proposed double-threshold behavior (6) that could provide an aggressive station with a significant increase in asymptotic utility. That is, besides leading to desirable per-stage strategy profiles, (6) also approximates equilibrium behavior for the repeated TRA game.

VII. Discussion

In this section we briefly discuss implementation issues concerning TRA detection, observation of relevant information by the network stations, and multi-hop considerations. It follows that the above described TRA discouragement is a challenging, yet technologically feasible scheme.

A. *Detecting Traffic Remapping*

The basis for detecting a TRA is determining if the monitored higher-layer traffic matches its class designation. In the case of EDCA, this means checking whether the structure of the user data contents in a data frame is characteristic of the traffic class that maps onto the AC specified in the MAC header (which, as noted before, is assumed not to have been tampered with). In particular, we need to recognize best-effort (e.g., HTTP, FTP, P2P) traffic being sent in the VO or VI ACs. This can be achieved with a variety of well-known traffic classification methods. These methods have thus far been mostly used either for attack detection, as part of an Intrusion Detection System (IDS), or for class-based QoS provisioning. Here, both these goals are addressed at once: we detect TRAs in order to provide appropriate QoS to each traffic class. The following well-known traffic classification methods, listed in order of increasing complexity, may be used, depending on permissible detection time and available resources:

- analysis of the transport protocol type and source and destination address/port is easy to implement, but may fail when non-standard ports or end-to-end encryption are used,
- deep packet inspection, i.e., performing pattern matching on the packet payload, available in both commercial products (Cisco's Network Based Application Recognition[8], Juniper's Application Identification[9], QOSMOS' Deep Packet Inspection[10]), as well as in open-source software[11,12], requires high processing power, and may also fail when non-standard protocols or end-to-end encryption are used, and
- statistical flow analysis based on machine learning [36], an emerging approach based on data-mining techniques such as payload length analysis, is efficient and allows the packet contents to be encrypted[13]. It does require a learning period but this can be performed offline.

[8] http://www.cisco.com/

[9] http://www.juniper.net/

[10] http://www.qosmos.com/

[11] http://l7-filter.sourceforge.net

[12] http://www.bro-ids.org

[13] Even with encrypted data, the tag which designates the traffic class (e.g., DSCP) must be transmitted in the clear.

This article has been accepted for publication in a future issue of this journal, but has not been fully edited. Content may change prior to final publication. Citation information: DOI 10.1109/TWC.2014.2321577, IEEE Transactions on Wireless Communications

Note that in contrast to typical traffic classification scenarios, flow monitoring for TRA detection need only be performed by selected stations and need not synchronize with the flows' lifetimes.

B. Observation of Relevant Information

In order to obtain relevant information for playing the repeated TRA game, the stations have to be able to observe:

- the throughput of their transmitted BE traffic,
- the QoS parameters (delay, jitter, PLR, PER) of their transmitted VO/VI traffic, and
- the DISSATISFACTION primitives from other stations.

Below we briefly address two implementation issues: how the above observations can be made and how much time they require in a nonstationary environment.

Measurement of end-to-end throughput is fairly straightforward and readily displayed in typical best-effort applications such as HTTP, FTP, or P2P. Such measurements are gathered using locally available information from TCP. QoS parameters can likewise be estimated in typical multimedia applications (such as VoIP) at the price of a certain processing overhead. These parameters can also be determined indirectly through degraded quality of experience (QoE) [37]. Finally, DISSATISFACTION primitives can be observed by any interested player that configures its wireless interface into promiscuous monitoring mode. Then, all incoming frames are analyzed, including those not destined for the player.

It would probably be difficult for a station to establish its satisfaction from the received service level on a timescale of less than one second. In our simulations we configured each stage of the TRA game to last 10 s. This was found enough to properly observe throughput and QoS parameters given that, when displayed by a measurement application, they are refreshed every several seconds. Also enough DISSATISFACTION primitives should appear within each stage if there is a dissatisfied station. A reduction of stage duration below 10 s (that is, more frequent strategy switching) would jeopardize the credibility of payoff values determined at the end of a stage, and the stations' boundedly rational, double-threshold play might not lead to high utilities.

C. Multi-hop Considerations

The consequences of a generalization of the proposed scheme to multi-hop networks should be studied at both the network and MAC layers. At the network layer, a TRA may backfire if the attacker interferes with a downstream node on the path to the attacker's destination. Alternatively, an attacker may attempt downgrading the priority of forwarded traffic [38]. This would add a whole new dimension to the problem and so is left out in the present study; yet it makes a very interesting avenue of future research. At the MAC layer, the presence of hidden stations calls for some extensions to the DISSATISFACTION scheme. The simplest is for an honest station, whether currently satisfied or dissatisfied, to relay DISSATISFACTION primitives whenever it senses them in its neighborhood. In a generic scenario, let an attacker station A and a dissatisfied honest station C be hidden from each other, but either be within the range of a station B. If B is honest then C acts upon A as it would in a single-hop network. If B is an attacker, the DISSATISFACTION primitives from C will ultimately force it to become honest and subsequently relay further DISSATISFACTION primitives.

VIII. Conclusions

Within the game-theoretic paradigm we have proposed a distributed scheme to discourage traffic remapping attacks (TRAs) in an ad hoc IEEE 802.11 EDCA environment. The scheme relies on broadcasting DISSATISFACTION primitives by honest stations (not attempting a TRA) currently dissatisfied with the network service level. These primitives are perceived by attacker stations as a threat of TRA detection and punishment, e.g., via higher-layer traffic classification and selective frame jamming, respectively, which they factor in their payoff functions in the arising TRA game. Contrary to some existing incentive compatible schemes, the DISSATISFACTION scheme does not require static identities of the stations or TTP services (which may be hard to come by in an ad hoc network). Additionally, it does not interfere with the IEEE 802.11 medium access function and has a low overhead thanks to passive monitoring (cf. Section VII.B) and piggybacking the DISSATISFACTION primitives.

In the presented repeated TRA game model, a station exhibiting boundedly rational, double-threshold behavior is able to learn if a TRA is necessary to reach satisfaction and at the same time harmless enough for the other stations to allow it. As a consequence, the stations either reach an 'all satisfied' strategy profile, or, if such a profile is not admitted by the traffic and network conditions, feel compelled to eventually retreat to a desirable 'all honest' strategy profile. Furthermore, the proposed strategy is either a subgame perfect equilibrium or, if not technically proven as such, there is little room for improvement of the asymptotic utility. Additionally, in the presence of aggressive stations (which cannot be satisfied while remaining honest and so perform TRAs persistently) the scheme tends to protect the payoffs of honest stations.

To determine whether given traffic conditions as well as EDCA configuration make a TRA setting requires further analytical studies, which we consider future work. Also, more complex scenarios of the TRA game need to be investigated. In particular, the proposed double-threshold behavior (6) would require modification for an increased number of ACs and perhaps redefining for more general class-based MAC protocols. Finally, the TRA game can be extended to multi-hop topologies, where on one hand the attack may fail because of multi-hop flow interactions and on the other the attack space increases in the form of downgrading forwarded traffic.

Appendix
Proofs of the Assertions

Assertion 1. *A stationary strategy profile is either 'all honest' or 'all satisfied' (if the latter is admitted by* $\mathbf{r}$*,* $\mathbf{t}$*, and* $\mathbf{d}$*), or both.*

Proof. Consider a stationary strategy profile $\hat{\mathbf{c}}$. From (6) it can be seen that if $\mathbf{c}^k = \hat{\mathbf{c}}$ then for $F_i^{m+1} = \{c_i^m\}$ to hold for all $m \geq k$, any BE station i must either be honest at $\hat{\mathbf{c}}$, with $u_i^m < \theta_{fi}$ for all $m \geq k$, or must observe $u_i^m \geq \theta_{xi}$ for all $m \geq k$. The former condition cannot occur since the nonnegative payoffs station i would receive after stage k would eventually drive u_i^m above θ_{fi} and result in arbitration between VO and BE, with VO ultimately chosen in some stage. Since $\mathbf{p}(\mathbf{c}^m)$ is constant for all $m \geq k$, the latter condition implies that all BE stations observe $p_i^m = 1$ for all $m \geq k$. Hence, $\hat{\mathbf{c}}$ is 'all satisfied' if all VO stations are satisfied; otherwise some VO station will broadcast DISSATISFACTION primitives and any BE station can only persist with $p_i^m = 1$ (i.e., ignore the DISSATISFACTION primitives) if it is honest, consequently $\hat{\mathbf{c}}$ is 'all honest'. □

Assertion 2. *Suppose the repeated TRA game is in stage k and that $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$ admit a strategy profile $\mathbf{c}^*$ with $\mathbf{p}(\mathbf{c}^*) = \mathbf{1}$. Then there exists an 'all satisfied' strategy profile (possibly different from $\mathbf{c}^*$) that is reachable from $\mathbf{c}^k$, i.e., one $\mathbf{c}^k$ may transform into after a finite number of stages under (6).*

Proof. We proceed by studying all possible cases in turn.

Case 1a: $\mathbf{p}(\mathbf{c}^k) = \mathbf{1}$, i.e., $\mathbf{c}^k$ is 'all satisfied'. The proof follows immediately.

Case 1b: $\mathbf{p}(\mathbf{c^k}) < \mathbf{1}$, i.e., there is a dissatisfied station at $\mathbf{c}^k$. This case breaks up into two subcases:

Case 2a: $\mathbf{c}^k = \mathbf{t}$, i.e., $\mathbf{c}^k$ is 'all honest'. Denote by A^* and H^* the sets of attacker and honest stations at $\mathbf{c}^*$, respectively. By Definition 1, part (b), $s_i(\mathbf{t}) = 1$ for all $i \in H^*$, therefore there can only be dissatisfied stations within A^* ($A^* = \emptyset$ leads back to Case 1a). Since by (6) a station can play honest arbitrarily long, receiving zero payoffs, it is feasible that u_j^m eventually falls between θ_{fj} and θ_{xj} for any BE station $j \in A^*$, making it free to play attacker. Again by Definition 1, part (b), the satisfaction of the stations in H^* will be then preserved. Hence, an 'all satisfied' strategy profile can eventually be reached.

Case 2b: $\mathbf{c}^k > \mathbf{t}$, i.e., there are attacker stations at $\mathbf{c}^k$. Then either $\mathbf{t}$ is reachable from $\mathbf{c}^k$, which leads back to case 2a, or not. The latter case implies that there is an attacker BE station i for which $u_i^m \geq \theta_{xi}$ (hence, c_i^m = VO) for all $m \geq k$. Again consider two subcases:

Case 3a: $s_j(\mathbf{c}^k) = 0$ for some station j that is honest at $\mathbf{c}^k$ ($c_j^k = t_j$). This is impossible, since if station j then plays honest long enough, its DISSATISFACTION primitives will cause persistent payoffs $p_i^m = -1$, which will eventually drive u_i^m below θ_{fi} and force station i to retreat to honesty.

Case 3b: $s_j(\mathbf{c}^k) = 1$ for all honest stations at $\mathbf{c}^k$, implying that some attacker stations are dissatisfied at $\mathbf{c}^k$ and receive zero payoffs. By a similar argument as above, one of those attacker stations eventually becomes free to switch to honest. If it then remains dissatisfied, we go back to Case 3a; otherwise we repeat our reasoning until the last dissatisfied attacker station switches to honest, whereupon we go back to Case 3a if this station remains dissatisfied, or to Case 1a if all the dissatisfied attacker stations have become satisfied. □

Assertion 3. *Suppose that $\mathbf{r}$, $\mathbf{t}$, and $\mathbf{d}$ admit an 'all satisfied' strategy profile. Then under (6), $\mathbf{u} \in \mathbf{1}^-$ is reachable from any stage of the repeated TRA game.*

Proof. By Assertion 2, an 'all satisfied' strategy profile is reachable. Let $\mathbf{p^k} = \mathbf{p}(\mathbf{c^k}) = \mathbf{1}$. Under (6), a BE station i for which $u_i^k \geq \theta_{fi}$ can maintain c_i^k indefinitely (having attained $u_i^k \geq \theta_{xi}$ after a finite number of stages). Consider a BE station j for which $u_j^k < \theta_{fj}$. Since $p_j^k = 1$, it must be that $u_j^{k-1} < u_j^k < \theta_{fj}$, hence c_j^k = BE, i.e., station j is honest at c^k. As such, it too can maintain c_j^k indefinitely. Therefore, $\mathbf{p^m} = \mathbf{1}$ for all $m \geq k$ is feasible, and the proof follows from (5). □

Assertion 4. *Suppose that (a) $N_{BE} < N$, (b) $\mathbf{p} = \mathbf{1}$ is not admitted, (c) BE station 1 is aggressive, and (d) some honest station different from station 1 cannot be satisfied in the presence of more than one attacker. Then under (6), $p_i(\mathbf{c}) \leq 0$ for all $\mathbf{c} \in \mathbf{F}$, i.e., an aggressive station cannot receive a positive payoff.*

Proof. By the definition of an aggressive station (cf. Section V.A), $s_1(\mathbf{t}) = 0$. Consequently, by Definition 1, part (b), we also have $s_1(t_1, \mathbf{c}_{-1}) = 0$ for any opponents' profile $\mathbf{c}_{-1}$ that involves any attackers. This implies that station 1 is never satisfied when it plays honest. Suppose now that station 1 is an attacker. If $\mathbf{c}_{-1}$ only consists of honest stations then either some of them are dissatisfied, in which case $p_1(\text{VO}, \mathbf{c}_{-1}) \leq 0$, or all are satisfied, in which case $\mathbf{p} = \mathbf{1}$ (contrary to assumption (b)) or $s_1(\text{VO}, \mathbf{c}_{-1}) = 0$. Thus the only possibility for station 1 to receive a payoff of 1 (i.e., be satisfied and not exposed) is to choose c_1 = VO against a $\mathbf{c}_{-1}$ that has the following properties: (i) $s_1(\text{VO}, \mathbf{c}_{-1}) = 1$, (ii) all the honest stations in $\mathbf{c}_{-1}$ (including a nonzero number of VO stations) are satisfied, and (iii) there is a dissatisfied station in $\mathbf{c}_{-1}$ (since $\mathbf{p} = \mathbf{1}$ is not admitted), which, in view of (ii), must be an attacker. Hence, (ii) and (iii) jointly contradict assumption (d). □

Acknowledgment

The work of J. Konorski is supported in part by the Ministry of Science and Higher Education, Poland, under Grant N519 581038, and by the Air Force Office of Scientific Research, Air Force Material Command, USAF, under Grant FA 8655-11-1-3076. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purpose notwithstanding any copyright notation thereon. The work of S. Szott is supported by the Polish National Science Centre (decision no. DEC-2011/01/D/ST7/05166).

References

[1] M. Natkaniec, K. Kosek-Szott, S. Szott, and G. Bianchi, "A Survey of Medium Access Mechanisms for Providing QoS in Ad-Hoc Networks," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 2, pp. 592–620, 2013.

[2] "IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, March 2012," 2012.

[3] G. Bianchi, A. Di Stefano, C. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental assessment of the backoff behavior of commercial IEEE 802.11b network cards," in *Proc. of INFOCOM*. IEEE, 2007, pp. 1181–1189.

[4] J. Konorski and S. Szott, "EDCA remapping in ad hoc IEEE 802.11 WLANs: An incentive compatible discouragement scheme," in *Wireless Days (WD), 2012 IFIP*, 2012. [Online]. Available: http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=6402868

[5] S. Szott and J. Konorski, "A game-theoretic approach to edca remapping attacks," in *Proc. of International Conference on Wireless Communications, Networking and Mobile Computing (WiCom)*, 2012.

[6] L. Galluccio, "A game-theoretic approach to prioritized transmission in wireless csma/ca networks," in *Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th*. IEEE, 2009, pp. 1–5.

[7] S. Szott, M. Natkaniec, and A. R. Pach, "An IEEE 802.11 EDCA model with support for analysing networks with misbehaving nodes," *EURASIP Journal on Wireless Communications and Networking*, vol. 2010, p. 71, 2010.

[8] P. Serrano, A. Banchs, and J. F. Kukielka, "Detection of malicious parameter configurations in 802.11 e edca," in *Global Telecommunications Conference, 2005. GLOBECOM'05. IEEE*, vol. 6. IEEE, 2005, pp. 5–pp.

[9] S. Szott, M. Natkaniec, and R. Canonico, "Detecting backoff misbehaviour in IEEE 802.11 EDCA," *European Transactions on Telecommunications*, vol. 22, no. 1, pp. 31–34, 2011.

[10] P. Kyasanur and N. H. Vaidya, "Selfish mac layer misbehavior in wireless networks," *Mobile Computing, IEEE Transactions on*, vol. 4, no. 5, pp. 502–516, 2005.

[11] M. Cagalj, S. Ganeriwal, I. Aad, and J.-P. Hubaux, "On selfish behavior in CSMA/CA networks," in *Proc. of INFOCOM*, 2005.

[12] L. Guang, C. Assi, and A. Benslimane, "Mac layer misbehavior in wireless networks: challenges and solutions," *Wireless Communications, IEEE*, vol. 15, no. 4, pp. 6–14, 2008.

[13] J. Konorski, "A game-theoretic study of csma/ca under a backoff attack," *IEEE/ACM Transactions on Networking (TON)*, vol. 14, no. 6, pp. 1167–1178, 2006.

[14] S. Szott, M. Natkaniec, and A. R. Pach, "Improving qos and security in wireless ad hoc networks by mitigating the impact of selfish behaviors: a game-theoretic approach," *Security and Communication Networks*, vol. 6, pp. 509–522, 2013. [Online]. Available: http://dx.doi.org/10.1002/sec.677

[15] K. Akkarajitsakul, E. Hossain, D. Niyato, and D. I. Kim, "Game theoretic approaches for multiple access in wireless networks: A survey," *Communications Surveys & Tutorials, IEEE*, vol. 13, no. 3, pp. 372–395, 2011.

[16] J. Hu, G. Min, W. Jia, and M. Woodward, "Admission control in IEEE 802.11e Wireless LAN: a game-theoretical approach," in *Game Theory for Wireless Communications and Networking*, Y. Zhang and M. Guizani, Eds. CRC Press, 2011.

[17] M. H. Cheung, A. H. Mohsenian-Rad, V. W. Wong, and R. Schober, "Random access protocols for wlans based on mechanism design," in *Communications, 2009. ICC'09. IEEE International Conference on*. IEEE, 2009, pp. 1–6.

[18] S. H. Nguyen, L. L. Andrew, and H. L. Vu, "Service differentiation without prioritization in ieee 802.11 wlans," in *Local Computer Networks (LCN), 2011 IEEE 36th Conference on*. IEEE, 2011, pp. 109–116.

[19] L. Zhao, L. Cong, H. Zhang, W. Ding, and J. Zhang, "Game-theoretic edca in ieee 802.11 e wlans," in *Vehicular Technology Conference, 2008. VTC 2008-Fall. IEEE 68th*. IEEE, 2008, pp. 1–5.

[20] M. Li and B. Prabhakaran, "Mac layer admission control and priority re-allocation for handling qos guarantees in non-cooperative wireless lans," *Springer Mobile Networks and Applications*, vol. 10, pp. 947–959, 2005. [Online]. Available: http://dx.doi.org/10.1007/s11036-005-4451-7

[21] P. Nuggehalli, M. Sarkar, and R. R. Rao, "Qos and selfish users: a mac layer perspective," in *Global Telecommunications Conference, 2007. GLOBECOM'07. IEEE*. IEEE, 2007, pp. 4719–4723.

[22] J. Price, P. Nuggehalli, and T. Javidi, "Incentive compatible mac-layer qos design," in *Proc. of 5th IEEE Consumer Communications and Networking Conference (CCNC)*, 2008.

[23] M. Ghazvini and N. M. K. Jamshidi, "Gtxop: A game theoretic approach for qos provisioning using transmission opportunity tuning," *PLoS ONE*, vol. 8, p. 8, 2013.

[24] S. Ross and B. Chaib-draa, "Satisfaction equilibrium: Achieving cooperation in incomplete information games," in *Advances in Artificial Intelligence*. Springer, 2006, pp. 61–72.

[25] ITU-T, "Recommendation Y.1541: Network performance objectives for IP-based services," December 2011.

[26] R. Stankiewicz, P. Cholda, and A. Jajszczyk, "QoX: What is it really?" *Communications Magazine, IEEE*, vol. 49, pp. 148–158, 2011.

[27] The Network Simulator NS-2. Online. [Online]. Available: http://nsnam.isi.edu/nsnam/index.php/Main_Page

[28] S. Wiethölter, M. Emmelmann, C. Hoene, and A. Wolisz, "TKN EDCA Model for ns-2," *Telecommunication Networks Group, Technische Universität Berlin, Tech. Rep. TKN-06-003*, 2006.

[29] E. Rasmusen, *Games and Information: An Introduction to Game Theory*. Blackwell Publishers, 2001.

[30] S. Szott, J. Gozdecki, K. Kosek-Szott, K. Loziak, M. Natkaniec, and I. Tinnirello, "The risks of wifi flexibility: Enabling and detecting cheating," in *Future Network and Mobile Summit (FutureNetworkSummit)*, 2013.

[31] A. Rubinstein, *Modeling bounded rationality*. The MIT Press, 1998, vol. 1.

[32] E. Friedman and S. Shenker, "Synchronous and asynchronous learning by responsive learning automata," *Unpublished manuscript*, 1996.

[33] D. Fudenberg and D. K. Levine, *The Theory of Learning in Games*. MIT Press, 1998.

[34] W. Feller, *An Introduction to Probability Theory and Its Applications*. J. Wiley and Sons, 1966.

[35] I. Marsic, "Computer networks: Performance and quality of service," *Rutgers University, New Jersey*, 2010.

[36] T. Nguyen and G. Armitage, "A survey of techniques for internet traffic classification using machine learning," *Communications Surveys Tutorials, IEEE*, vol. 10, no. 4, pp. 56–76, 2008.

[37] M. Fiedler, T. Hossfeld, and P. Tran-Gia, "A generic quantitative relationship between quality of experience and quality of service," *Network, IEEE*, vol. 24, pp. 36–41, 2010.

[38] S. Szott, M. Natkaniec, and A. Banchs, "Impact of Misbehaviour on QoS in Wireless Mesh Networks," in *Proc. of IFIP Networking*, 2009.