On the Exploitation of CDF based Wireless Scheduling

Udi Ben-Porat[a,*], Anat Bremler-Barr[b], Hanoch Levy[c]

[a]Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Switzerland.
[b]Computer Science Dpt., Interdisciplinary Center, Herzliya, Israel.
[c]Computer Science Dpt., Tel-Aviv University, Tel-Aviv, Israel.

Abstract

Channel-aware scheduling strategies - such as the CDF scheduler (CS) algorithm - provide an effective mechanism for utilizing the channel data rate for improving throughput performance in wireless data networks by exploiting channel fluctuations. A highly desired property of such a scheduling strategy is that its algorithm is stable, in the sense that no user has incentive "cheating" the algorithm in order to increase his/hers channel share (on the account of others). Considering a *single user* we show that no such user can increase his/hers channel share by misreporting the channel capacity. In contrast, considering a *group of users*, we present a scheme by which coordination allows them to *gain permanent increase* in both their time slots share and in their throughput on the expense of others, by misreporting their rates. We show that for large populations consisting of regular and coordinated users in equal numbers, the ratio of allocated time slots between a coordinated and a regular user converges to $e - 1 \approx 1.7$. Our scheme targets the very fundamental principle of CS (as opposed to just attacking implementation aspects), which bases its scheduling decisions on the Cumulative Distribution Function (CDF) of the channel rates reported by users. Our scheme works both for the continuous channel spectrum and the discrete channel spectrum versions of the problem. Finally, we outline a modified CDF scheduler immune to such attacks.

Keywords: Wireless, Cellular, MAC layer, CDF, Scheduling, Fairness, DDoS, Attack, Exploit.

1. Introduction

High-speed wireless networks are becoming increasingly common and along with that the strategy of *scheduling* the high-speed data - which is vital to the performance of moderns wireless systems - has become the subject of active research. The modern wireless networks standards such as HSPDA [1] and EV-DO [2] [3] allow new generation of channel aware schedulers - such as the Proportional Fairness [4] [5] and the CDF scheduler [6] - which improve throughput performance by exploiting channel fluctuations while maintaining fairness between the users.

The CDF Scheduler (CS) makes scheduling decisions based on the Cumulative Distribution Function (CDF) functions of the users in such way that every time slot the user whose rate is the least probable to become higher is scheduled for transmission. An important property of this scheduler is that it statistically allocates all users an equal number of slots while smartly utilizing the knowledge of channel capacity to dynamically select at every moment the more attractive (higher capacity) users. A distinctive feature of this algorithm is that it allows to predict the exact throughput for each user based on his/hers[1] CDF alone, regardless of changes in the channel rate distribution of other users. These features and its simple notion of fairness (equal time share) make CS an attractive alternative to the Proportional Fairness Scheduler (PFS) [4]. Recent studies [7] [8] revealed the vulnerability of PFS to delays/jitter and loss of throughput caused by malicious users by providing false channel capacity reports. In this paper the vulnerability of the CDF scheduler to threats of non-conforming opportunistic users as well as malicious users is investigated for the first time.

One of the main roles of a resource allocation mechanism is to ensure fairness of the allocation under the assumption that every user aims at increasing his own allocation. Furthermore, it is highly important that the

*Corresponding Author.
E-mail addresses: ehudb@tik.ee.ethz.ch (U. Ben-Porat), bremler@idc.ac.il (A. Bremler-Barr) and hanoch@post.tau.ac.il (H. Levy).

[1]From now on we use "he" and "his" to mean "she/he" and "hers/his" for the sake of reading flow.

scheduler will be resilient to users who may try to increase the resources allocated to them by not fully conforming with the protocol rules.

The objective of this work is to study this problem. Namely, whether a user, or a group of users, can mislead the CDF scheduler by providing false channel capacity reports and use it to increase the amount of resources allocated to them. Every modern channel-aware scheduler must allow a temporary state of unfairness in order to utilize a temporary exceptionally good channel condition of one of the users. Nevertheless, it is still expected that in the long run - that is, in the steady state - fairness is enforced. For example, in [7] the authors presented an attack on PFS in which a starved user can suddenly report an exceptionally good channel condition and temporarily be granted high priority, which cause other users to experience jitter. However, in the long run, the fairness that PFS is meant to ensure, is kept. In this work, we show that the CDF scheduler can be attacked by malicious and selfish users who gain a *permanent* advantage over users. That is, the time share fairness that the CDF scheduler is meant to ensure is not kept even in the steady state. We show that this is a fundamental weak point of the CDF scheduler regardless of its exact implementation.

To this end we show that the CDF algorithm is resilient against "attacks" produced by a *single user*. That is, a single user can increase neither the number of slots nor the bandwidth allocated to him by providing misleading information about his channel capacity. We then show, that nonetheless, a *group of coordinated users* which collaborate with each other can *increase both* the *number of slots* and the *bandwidth* allocated to *each of them*. That is, while the scheduler is designed to counter an independent selfish behavior of a single user, its design does not take into account the possibility of a coordinated group of users. The capacity announcement strategy used by the coordinated users is very simple and requires only knowledge of each other's capacity. We conduct the analysis of this strategy and derive its performance gains. The analysis is carried out both for the the continuous rate distribution model (Section 3) and the discrete rate distribution model (Section 4). Our results show that the gain that such non-conforming users can achieve may be as high as 28% in a typical system configuration (30 users). Furthermore, the ratio between the slot allocation of a coordinated user and a regular user can reach $e - 1 \approx 1.7$. We further consider coordinated *malicious* users. These aim at reducing the performance the regular users, not caring about their own performance. We show that the channel share loss that the regular innocent users suffer can be as high as 48% in a typical system configuration.

The attack algorithm we show exploits the stochastic worst case traffic pattern of multiple users that can be applied to the system. This type of attack is demonstrated in the Reduction of Quality (RoQ) attacks papers [9, 10, 11]. RoQ attacks target the adaptation mechanisms by hindering the adaptive component from converging to steady-state. This is done by sending - from time to time - a very short burst of surge demand imitating many users and thus pushing the system into an overload condition. Using a similar technique, Kuzmanovic and Knightly [12] presented the Shrew Attack which is tailored and designed to exploit TCP's deterministic retransmission timeout mechanism. Another example of an attack exploiting the stochastic worst case is given in [13, 14]. There it is shown that Weighted Fair Queueing (WFQ), a commonly deployed mechanism to protect traffic from DDoS attacks, is ineffective in an environment consisting of bursting applications such as the Web client application. The paper [15] shows attack on the SSL handshake, by requesting again and again hard SSL requests.

The rest of the paper is organized as follows: After model and preliminaries given in Section 2, Section 3 analyzes non-conformist users under the continuous rate distribution, and Section 4 does it under the discrete rate distribution. In Section 5 we analyze the loss for regular users by coordinated and malicious users in the practical discrete model. Finally, in Section 6 we outline a modified CDF Scheduler immune to selfish or malicious behavior. Note that a short abstract of this work has been presented at [16].

2. Assumptions, Model and Preliminaries

In the scheduling models discussed in this work, time is slotted to slots $t = 1, 2, ...$ and the possible channel rates are arbitrary and non negative. The rate at which user k can transmit at time slot t is given by $R_k(t)$. $R_k(t)$ is distributed according to random variable R_k associated with user k, and whose CDF is $F_{R_k}(r) = Pr[R_k \leq r]$. $R_k(t)$ is a stationary random process assumed to be independent of $R_k(t')$ for any $t \neq t'$ and of $R_j(t')$ for any $j \neq k$ and any t'.

At each slot t, each user k announces to the scheduler his actual value $R_k(t)$. The scheduler may compute the distribution $F_{R_k}(r)$ from the past reports of user k. Note that we demonstrate the vulnerability of CS without targeting a weak point in the inferring mechanism, therefore throughout the paper we assume the schedular has the precise CDF functions of the channel rates reported by users. At time t, the scheduler can use both the studied $F_{R_k}(r)$, $k = 1, ..., K$ and the current user rates $R_k(t)$, $k = 1, ..., K$ to decide to which user to transmit at slot t. The rate at which the server will transmit to the selected user, say k, is $R_k(t)$.

3. The Basic Problem: Dealing with Continuous Rate Distributions

In this section we assume that all the channel rate distributions are continuous, that is the distribution functions do not contain mass values (i.e. $F_{R_k}(r)$ is differentiable and $Pr[R_k = x]$ equals zero for every x). Later, in Section 4, we will deal with discrete (and mixed) probability functions.

3.1. Scheduling Algorithm

The basic CDF Scheduler (CS), aiming at dealing with continuous distributions, was introduced in [6] and operates as follows. Recall that $R_k(t)$ is the actual channel capacity of user k at time slot (TS) t and let $k^*(t)$ be the user selected for data transmission. The scheduler selects $k^*(t)$ to be the user for which $P[R_k > R_k(t)]$ is the smallest among all users. That is, the user whose rate is the least probable to become higher, namely:

$$k^*(t) = argmax_k\{F_{R_k}(R_k(t))\},$$

where $F_{R_k}(r) = P[R_k \leq r]$. The original scheduler definition [6] includes the option to assign each user a special weight w_k according to $k^*(t) = argmax_k\{F_{R_k}(R_k(t))^{\frac{1}{w_k}}\}$ but for the sake of simplicity we omit the weight factor and assume the Base Station (BS) serves the users equally. For notational simplicity we define $V_k(t) = F_{R_k}(R_k(t))$ and since $k^*(t) = argmax_k\{V_k(t)\}$ we will refer to $V_k(t)$ as the priority value assigned to User k at TS t.

The CDF scheduler relies on a well-known property of CDF functions to ensure time share fairness: The CDF function of every continuous random variable X is distributed uniformly, $F_X \sim \text{Uniform}(0, 1)$. That is, the priority value of every user k is distributed uniformly, $V_k \sim \text{Uniform}(0, 1)$, regardless of the distribution of R_k. Therefore, all users have equal chance to obtain the highest priority value and hence time-share fairness is maintained.

3.2. Misreporting of channel rates cannot benefit a single user

The idea of users reporting fake channel rate to exploit the properties of channel aware scheduler was already introduced in [8]. The users can fake channel rate by modifying their laptop' 3G PC cards, by either through the accompanying software development kit or the device firmware. And the providers cannot detect it, even if they attempt tamper-proof technique [8]. Nevertheless, we prove that under the CDF scheduler, it is impossible for one user to benefit from an additional time share or throughput by misreporting his channel condition. We show that this result stems from the fundamental characteristics of the CDF scheduler which are common to both the continuous and the discrete models (Section 4).

Theorem 1. *If a user has no knowledge on the reports of others, then no strategy can benefit him with long-run additional time share.*

Proof: The CDF function of every continuous random variable X is distributed uniformly, $F_X \sim \text{Uniform}(0, 1)$, regardless of its distribution. Therefore, since the channel condition reports issued by the users are independent of each other, the priority value of every user is a random number between zero and one. Hence, every user has an equal probability to win a time slot. Therefore, all users receive (in the steady state) the same amount of time share $\frac{1}{N}$ where N is the number of users in the system. □

Note that the proof for Theorem 1 does not hold for a user who coordinate with other user(s), because then the channel reports of the users are no longer independent. In addition to Theorem 1, under both continuous and discrete models, there is no misreporting strategy that allows a user to achieve higher *throughput* than what he would get by always reporting his real rate. Due to its length, the formal proof of this claim (under the discrete model, which is the model used in practice) is placed at the Appendix. Nevertheless, we give here an intuitive explanation for the validity of this claim (under both continuous and discrete models). A false reports strategy can involve reporting a fake high rate when the real rate is lower and vice versa. Reporting a fake high channel rate might increase the priority values in some of the time slots in which the real rate is low. However, it makes the high rates reported by the user less exceptional than they really are and hence decrease the priority value when the user truly experiences and reports a high rate. Such behavior is not beneficial since increasing the priority value of some low-rate time slots on the expense of the priority value of high-rate time slots, never benefits the user. Reporting a rate lower than the real rate has also both negative and positive effects. The positive effect is that such behavior makes other time slots in which the user report the truth about his high rate look more exceptional. Hence, the user will have an increased priority value in these time slots. The negative effect is that in the time slots in which he fakes a low rate, the user has a lower priority value than if he reported the truth. In addition, even if he is assigned for transmission, the transmission rate will fit his report and will be lower than what his real channel condition can support. In the formal proof we show that this negative effect shadows the positive effect. Hence, if to summarize it in one sentence, scarifying some of the high rate slots to increase the priority value in the rest never pays off. Finally we can conclude that since no strategy can benefit the user with either additional throughput or time share – the best strategy for one user is to always report his real rate.

3.3. Coordinated Users Strategy

We next deal with a group of non-conforming opportunistic users who coordinate their action and reporting in order to increase their time and throughput shares.

In the previous section we explained that the negative effects of a misreporting strategy exceed its benefit when one user acts on his own. In this section we describe a cooperation scheme in which a group of users can gain additional time share and throughput while avoiding the negative effects that reporting a false channel condition may cause.

Let C be a group of $|C| = L$ coordinated users and N is the number of additional regular users in the network. Each one of the coordinated users knows if his rate is the least probable to be higher (and therefore will get the highest scheduling priority) in C before reporting to the BS. Let $c^*(t)$ be the user with the highest CDF value in the group in time t. Formally, $c^*(t) = argmax_c\{F_{R_c}(R_c(t))\}$ where $c \in C$. The reporting strategy is simple: at time slot t user c^* will be the only one acting normal (reporting his real channel rate $R_{c^*}(t)$) while all others report zero.

The users share their CDF value, so each user knows if he is c^* or not. The coordinated users strategy can be implemented using a low bandwidth medium/side-channel that allows the users to share this small amount of information. For example, a big factory using a designated private (low rate and cost) wireless network to coordinate the access points used by its employees in order to gain more throughput to its users. Note that in Section 5 we describe a malicious strategy that does not require any communication between the users during the attack and causes even a greater damage to the system than the coordinated strategy described here.

3.4. Analysis of the Coordinated Users Share

Let R'_c be the R.V. of the *reported* channel rate by user c when he follows the coordination strategy. Recall that $R_c(t)$ is the *real* channel rate of user c at TS t, therefore $R'_c = R_c$ if the user behaves normally.

Lemma 1.

$$F_{R'_c}(r) = G(F_{R_c}(r)) \tag{1}$$

where $G(x) = \frac{L-1}{L} + \frac{1}{L}x^L$ *and* $c \in C$.

Proof: Let E_r be the event where user $c \in C$ reports to the BS channel rate less than r ($R_c(t) < r$). Let WIN be the event where user c is the chosen user in C to report his real channel rate, therefore:

$$F_{R'_c}(r) = (1 - P[WIN])P[E_r|\neg WIN] + P[WIN]P[E_r|WIN] \tag{2}$$

Every time slot each coordinated user has an equal probability to be the one reporting his real channel rate, therefore $P[WIN] = \frac{1}{L}$. If c is not chosen to report his real rate, then he reports minimal channel rate, therefore for every r, $P[E_r|\neg WIN] = 1$. Now all is left is compute $P[E_r|WIN]$.

Let R.V. Y be $Y(t) = max_j\{F_{R_j}(R_j(t))|j \in C\}$, since F_{R_j} is CDF then $P[F_{R_j} < x] = x$ and we get $P[Y < y] = y^L$, this is true regardless of the ID of j for whom $F_{R_j}(R_j(t)) = Y(t)$, therefore $P[Y < y] = P[Y < y|WIN] = y^L$. According to the strategy, if c reports less than r and he is the chosen user (Event $E_r \bigwedge WIN$) then his real channel rate – $R_c(t)$ – must be less than r which means that $F_{R_c}(R_c(t)) < F_{R_c}(r)$. Given WIN, then $F_{R_c}(R_c(t)) = max_j\{F_{R_j}(R_j(t))|j \in C\} = Y(t)$ and according to what we just showed we get $(E_r \bigwedge WIN) \iff F_{R_c}(R_c(t)) < F_{R_c}(r) \iff (Y(t) < F_{R_c}(r))$ and we get $P[E_r|WIN] = P[Y(t) < F_{R_c}(r)] = (F_{R_c}(r))^L$ since $P[Y < y|WIN] = y^L$ as we showed earlier. According to Equation 2 and the conditional results we showed we get Equation 1. □

Recall that $V_c(r) = F_{R_c}(r)$ is the priority value of user c for reporting channel rate r to the BS when he behaves normally (always reporting his real channel rate). Now let $V'_c(r) = F_{R'_c}(r)$ be the priority value he gets for r when following the coordination strategy, then according to Lemma 1 we get that:

$$V'_c(r) = G(V_c(r)) \tag{3}$$

Before we analyze the benefit from coordination, we first prove that it can never harm (and hence only benefit) the coordinating users.

Theorem 2. *User following the coordination strategy will still win every time slot that he would have won when if he behaved normally. Therefore, his throughput and time share can only be increased when following the coordination strategy.*

Proof: Assume user c_0 obtains the highest priority value when all the users in the system behave normally. Therefore, $c_0 = k^*(t) = argmax_k\{V_k(R_k(t))\}$. Since also $c_0 = c^*(t)$, when following the strategy, c_0 will be the user from C reporting his real rate in this time slot. A simple function analysis can show that $G(x) > x$ for $x \in (0, 1)$. Therefore, according to Lemma 1, the priority value he gets $V'_{c_0}(r) = G(V_{c_0}(r))$ is greater than $V_{c_0}(r)$ which is greater than the priority values of all other users. Therefore, he still obtains this time slot when following the strategy. □

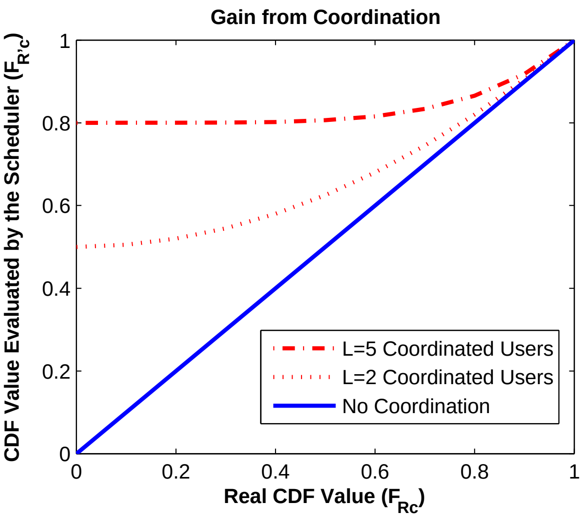

Figure 1: The Y-axis is the CDF value evaluated by the base station ($F_{R'_c}(r)$) as a function of the real CDF values that would have been evaluated without coordinating with other users ($F_{R_c}(r)$).

Fig. 1 depicts the gain from following the coordination scheme based on Eq. 1. For example, let r_0 be a rate with a corresponding CDF value of $F_{R_c}(r_0) = 0.4$. That is, in reality, 40% of the time the channel rate of the user is r_0 or lower. For normal users, who always report their real rates, the scheduler evaluates their real CDF value $F_{R'_c}(r) = 0.4$ as the solid curve (labeled *No Coordination*) shows. If the user coordinates with another user, $L = 2$, then, as the dotted curve show, the scheduler evaluates a CDF value of 0.58 instead of 0.4. As seen in the dashed curve, $L = 5$, when he cooperates with four other users, his priority value will be 0.802 instead of 0.4. As explained in the proof of Lemma 1, when a user coordinates with $L - 1$ other users, he is expected to report a zero channel rate $L - 1/L$ of the time. Therefore, when he finally reports a non-zero rate $r > 0$, then $F_{R'_c}(r) = Prob(R'_c < r) > L-1/L$. This explains why the values of the non-solid curves in Fig. 1 are greater than $L-1/L$. For example, assume a system with three users, U_1, U_2 and U_3. If U_1, who always report his true rate, has a CDF value of 0.4, what is the probability that he will be assigned with the time slot? In the case where U_2 and U_3 are regular (non-coordinated) users, the probability that both of them have priority value that does not exceed 0.4 is $0.4^2 = 16\%$. However, if U_2 and U_3 are coordinated, then one of them is going to report his real rate, which will be given with a CDF value higher than 0.5 regardless of the rate he reports. Therefore, if U_2 and U_3 are coordinated,the probability of U_1 to win the time slot with a CDF value of 0.4 is 0% instead of 16%. To conclude, the above results show how users in a coordinated group systematically increase their priority values. In the following theorems we show how this advantage translates into a larger time share and more throughput.

Theorem 3. *In a network consisting $L + N$ users of which L are coordinated, the time share fraction dedicated to the L coordinated users (jointly) depends only on L and N (regardless of the channel rate distributions of any of users) and is given by:*

$$\frac{L}{N+1}\left(1-\left(\frac{L-1}{L}\right)^{N+1}\right). \tag{4}$$

Proof: First we make few definitions and short calculations to be used later. Let R.V. W be the maximal priority $W(t) = max_c\{V_c(R_c(t))\}$. Since $P[V_c < x] = x$, we get $P[W < w] = w^L$. Therefore, the probability density function (PDF) of W is $f_W(w) = (P[W < w])' = L{\cdot}w^{L-1}$.

Let R.V. $B(t) = max_n\{V_n(R_n(t))\}$, where n is one of the regular users, be the highest priority among the regular users and R.V. $A = max_n\{V'_n(R_n(t))\}$ be the highest priority among the coordinated users. Then, from Equation 3 we get that $A = max_n\{G(V_n(R_n(t)))\}$. In addition, since G is monotonically increasing, then $max_n\{G(V_n(R_n(t)))\} = G(max_n\{V_n(R_n(t))\})$ and we get $A = G(W)$.

According to that we get:

$$\begin{aligned} P[B < A] &= \int_{w=0}^{1} f_W(w) \cdot P[B < A|W = w]dw \\ &= \int_{w=0}^{1} L \cdot w^{L-1} \cdot P[B < G(w)]dw. \end{aligned}$$

The probability that all regular users will have priority less than a is $P[B < a] = a^N$, then $P[B < G(w)] =$

$(G(w))^N$ and we get:

$$\begin{aligned} P[B < A] &= \int_{w=0}^{1} L w^{L-1} \left(\frac{L-1}{L} + \frac{1}{L} w^L \right)^N dw \\ &= \int_{w=0}^{1} L w^{L-1} \cdot \sum_{i=0}^{N} \binom{N}{i} \frac{(L-1)^{N-i}}{L^N} w^{Li} dw \\ &= \sum_{i=0}^{N} \binom{N}{i} \cdot \frac{(L-1)^{N-i}}{L^{N-1}} \int_{w=0}^{1} w^{(i+1)L-1} dw \end{aligned}$$

and since $\int_{w=0}^{1} w^{(i+1)L-1} dw = ((i+1)L)^{-1}$ we get:

$$P[B < A] = \sum_{i=0}^{N} \binom{N}{i} \frac{(L-1)^{N-i}}{L^N (i+1)}.$$

Hence, we receive that

$$P[B < A] = \frac{L}{N+1} \left(1 - \left(\frac{L-1}{L} \right)^{N+1} \right) \quad (5)$$

□

This result now allows us to evaluate the "inequality" in time slot allocation between a coordinated user and a regular user. This "inequality" can be evaluated by the ratio between the slot shares of these users, which as shown next, can be very high:

Corollary 4. *Let C^{share} and R^{share} be the time share of the coordinated users and regular users respectively. When $N = L-1$, then $\lim_{L\to\infty} C^{share} = 1 - e^{-1}$ and $\lim_{L\to\infty} R^{share} = e^{-1}$. This means that $\lim_{L\to\infty} \frac{C^{share}}{R^{share}} = e - 1 \approx 171\%$ although this ratio - which equals $L/N = L/(L-1)$ under normal conditions - should converge to 100%.*

Theorem 5. *In a network consisting $L+N$ (in the continuous model) users of which L are coordinated, the average throughput (per time slot) of a coordinated user c is given by:*

$$\sum_{i=0}^{N} \binom{N}{i} \cdot \frac{(L-1)^{N-i}}{L^N} \int_{w=0}^{1} w^{(i+1)L-1} \cdot F_{R_c}^{-1}(w) dw, \quad (6)$$

where $F_{R_c}^{-1}(w)$ is the inverse function of F_{R_c}, the CDF of the real channel rate distribution.

Proof: The proof is very similar to the proof of Theorem 3, we will skip the identical parts in the proof. Assume A, B and W as defined in the proof of Theorem 3. Let R.V. D be the rate that user c receives (at some time slot t). If $W = w$ then user c gets throughput from the BS only if $B > A$ (this is a time slot which is obtained by a user from C). In addition, c has to be the user chosen among the coordinated users to transmit his real channel rate (happens with probability $\frac{1}{L}$). When $W = w$ and c is the chosen user, then $F_{R_c}(r) = w$ therefore his the rate in this time slot is given by $F_{R_c}^{-1}(w)$, therefore we get:

$$\int_{w=0}^{1} f_W(w) \cdot P[B < A | W = w] \cdot \frac{1}{L} F_{R_c}^{-1}(w) dw.$$

Continuing the calculations in the same manner as in the proof of Theorem 3 until the last equation which consisting the integral, will give us Equation 6. □

3.5. Evaluation and Discussion

In Figure 2 we evaluate the relative benefit in time slots a coordinated user gains from the coordination strategy; this is relatively to what he would get ($\frac{1}{N+L}$) if he did not coordinate. The figure depicts this relative benefit as a function of the number of coordinated users L (given that the total number of users is fixed $N + L = 30$). One may observe that the relative benefit is maximized at $L = 11$, implying that a coalition of 11 coordinated users has only a little incentive adding more users to the coalition. The relative benefit per user obtained is 28%.

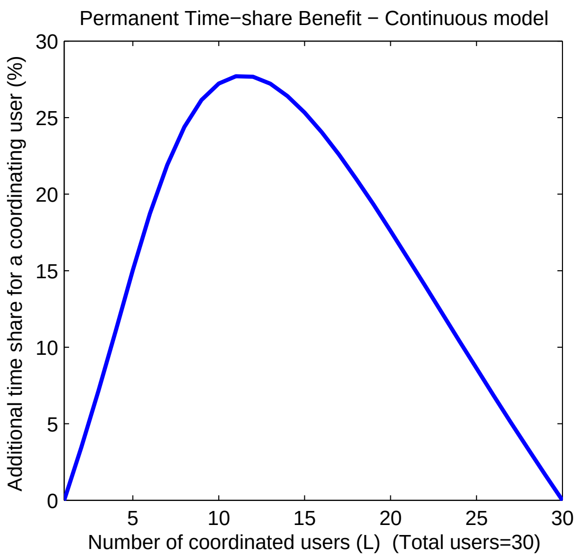

Figure 2: Y-axis is the additional time share (in percents) that a coordinated user gains when he takes part in a coordinated group of size L (X-Axis). The results show that when there are 30 users in the system, participating in a coordinated group of 11 users is the most beneficial and increases the time share of the user by 28%.

When $L = N + 1$ - which happens approximately at $L = 15$ in Figure 2 - the time share of the L users when behaving normally is given by $\frac{L}{2L+1}$ and converges to $\frac{1}{2}$. Corollary 4 pointed out that when $L = N + 1$ the time share obtained by the coordinated users converges to $1 - e^{-1}$ instead of $\frac{1}{2}$, means a coordinated user benefits from an additional time share of $(1 - e^{-1})/(\frac{1}{2}) - 1 = 26.4\%$ which is close to the result at $L = 15$ which equals 25.3%. Recall that according to Theorem 3 these results are valid for every system with 30 users regardless of the channel rate distributions of the users.

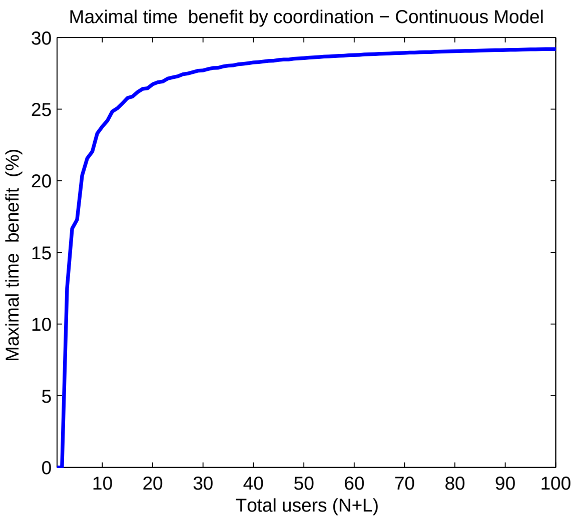

Figure 3: Maximal time share (percents) that one user benefits from following the coordination strategy as a function of the total population in the system (X-Axis).

In Figure 3 we evaluate this maximal benefit as a function of the population size ($N + L$ varies). One may see that this maximal benefit per user (of the coalition) monotonically increases with the population size and approaches 30% at large populations.

According to Theorem 2 the throughput of the coordinated users can only be increased. The throughput result for each user can be calculated according to Theorem 5 and depends on the user's specific distribution of channel rates and can be different for different users in the coordinated group. The throughput gain for a coordinated group is given in the evaluation for the discrete model.

4. CS with discrete channel rates range

4.1. Scheduling Algorithm

The original version of the CS algorithm ([6]) assumed continuous channel rate values even though practical systems use discrete values. We now summarize the extension of the CS algorithm to the case of discrete channel rate values which appears in [17]. Again, to keep the calculations simple, we assume all users have the same weight and exclude the weight factor.

In the discrete model $R_k(t) \in \{r_1, r_2, ..., r_M\}$ where $r_1 < r_2 < ... < r_M$. At TS #t user k feeds back $m_k(t) \in \{1, ..., M\}$ the index of his channel rate value. Denote $q_{k,m} \equiv F_{R_k}(r_m) = \sum_{i=1}^{m} P[R_k = r_i]$ where $q_{k,0}$ is set to 0 for notational convenience. Instead of simply taking $q_{k,m_k(t)}$ to be the priority value of user k, the CDF scheduler generates for each user a random priority given by a R.V. $U_k(t)$ which is uniformly distributed in the interval $[q_{k,m_k(t)-1}, q_{k,m_k(t)})]$. Finally, the scheduler selects the user with the highest priority $k^*(t) = argmax_k\{U_k(t)\}$. The priority value of the discrete range algorithm (U_k) preserves the fundamental characteristic of the priority value of the continuous range algorithm (V_k) which is that for every user k we get $P[U_k(t) \leq x] = P[V_k(t) \leq x] = x$. More precisely, as in the continuous model, the priority value of every user is distributed uniformly in $[0, 1]$ (regardless of his channel rate distribution) and hence time share fairness maintained by the scheduler also under the discrete model.

4.2. Misreporting of channel rates cannot benefit a single user

Section 3.2 explains why a single user cannot benefit neither additional time share nor throughput by misreporting his channel rate on his own. Both under the discrete and the continuous model, the priority values of all users are distributed uniformly in $[0, 1]$. Therefore, Theorem 1 which proves that one user cannot gain additional *time share* is valid also under the discrete model. To complete the proof, we provide at the Appendix a formal proof under the discrete model for the claim that no strategy can benefit a single user who acts alone with additional *throughput*.

4.3. Coordinated Users Strategy

Assume a coordinated group of users C ($|C| = L$) with the same channel rate probability (but independent from each other), means $\forall_{c_1,c_2 \in C}\ q_{c_1,i} = q_{c_2,i}$. Let $m^*_c(t) = argmax_i\{r_i | \exists c.R_c(t) = r_i\}$ where c are users from the coordinated group. Every time slot, only the users with channel rate $r_{m^*_c(t)}$ will report their real channel rate while all other report the lowest channel rate possible r_1 (they have no chance getting the highest priority).

Let $p_{c,i}$ be the probability that a coordinated user *experiences* channel rate r_i and $p'_{c,i}$ be the probability that he *reports* channel rate r_i. Then by following the strategy we get:

Corollary 6. *From the point of view of the BS (Base Station), the channel rate probability of coordinated user c is given by p' as follows:*

$$\begin{aligned} p'_{c,1} &= 1 - \sum_{i=2}^{M} p'_{c,i} \\ p'_{c,i} &= p_{c,i} \cdot (q_{c,i})^{L-1} \quad {\scriptstyle (i \geq 2)} \end{aligned}$$

Then when a coordinated user c reports $m_c(t) = i$, the scheduler generates a uniform priority value in the interval $[q'_{c,i}, q'_{c,i-1}]$ where $q'_{c,i}$ is the CDF of the reported channel rates of user c - $q'_{c,i} = \sum_{j=1}^{i} p'_{c,j}$.

For example, assume a user c_1 in a network with $p_{c_1,1} = p_{c_1,2} = p_{c_1,3} = \frac{1}{3}$ to have one of the three possible channel rates r_1, r_2, r_3. In a normal situation, if user c_1 reports channel rate r_1, then the CS generates a priority value U_{c_1} in $[0, \frac{1}{3}]$, if he reports r_2 then the range is $[\frac{1}{3}, \frac{2}{3}]$ and finally if he reports r_3 the range will be $[\frac{2}{3}, 1]$. Now assume this user is part of a coordinated group (where all users share the same channel rate distribution) which follows the coordinated users' strategy and we want to find $p'_{c_1,i}$ which is the probability for him to actually report rate r_i. When he experiences r_3 then he will always report r_3 because there is no other coordinated user who will surely have higher. If he coordinates with two more users who have the same rate probabilities then according to Corollary 6 we get that when c_1 reports r_2 the CS will generate U_{c_1} in the interval $[\frac{14}{27}, \frac{2}{3}]$ and this increases his expected U value. Therefore, the expected number of time slots where he obtains the highest priority.

4.4. Analysis of Coordinated Users Share

Lemma 2. *Let $MAXj, i$ be the event where $r_{m_c^*(t)} = r_i$ is reported by exactly j of the coordinated users.*

$$\begin{aligned} P[MAX_{j,1}] &= 0 \quad {\scriptstyle (\forall j<L)} \\ P[MAX_{L,1}] &= (p_{c,1})^L \\ P[MAX_{j,i}] &= \binom{L}{j} (p_{c,i})^j \cdot (q_{c,i-1})^{L-j} \quad {\scriptstyle (\forall i \geq 2)} \end{aligned}$$

Proof: $r_{m_c^*(t)} = r_i$ means that r_i is the highest rate of the users in C, therefore $i = 1$ only when all the users in C have r_1 which happens with probability of $(p_{c,1})^L$. It is impossible that $j < L$ coordinated users will report r_1 since it means that at least one user reports higher channel rate than r_1 which contradicts $r_{m_c^*(t)} = r_1$ and therefore $P[MAX_{j<L,1}] = 0$. There are $\binom{L}{j}$ possible combination of j users in C and each such combination reports $r_{m_c^*(t)} = r_i$, ${\scriptstyle (i \geq 2)}$ only when their actual channel rate is r_i (with probability $(p_{c,i})^j$) and the channel rate of all others is r_{i-1} or less (with probability $(q_{c,i-1})^{L-j}$). □

Lemma 3. *Let C_{WIN} be the event where some coordinated user wins a time slot, then*

$$\begin{aligned} P[C_{WIN}|MAX_{j,1}] &= (q'_{c,1})^N \cdot \frac{j}{j+N} \\ P[C_{WIN}|MAX_{j,i}] &= \sum_{s=0}^{N} \binom{N}{s} (p'_{c,i})^s (q'_{c,i-1})^{N-s} \cdot \frac{j}{j+s} \quad {\scriptstyle (i \geq 2)} \end{aligned}$$

Proof: According to $P[U_k(t) \leq u] = u$, the probability for regular users to get a priority value u is uniform in $[0, 1]$. When $m_c^*(t) = i \geq 2$, a coordinated user can obtain the highest priority value only when all N users has priority values in $[0, q'_{c,i}]$. With probability of $\binom{N}{s}(p'_{c,i})^s(q'_{c,i-1})^{N-s}$ exactly s regular users have priority values in $[q'_{c,i-1}, q'_{c,i}]$ while all the other $N - s$ regular users have no chance getting the highest priority. Given that event, the priority values - U_s of each of the s regular users and U_c of the coordinated users with r_i is uniform in $[q'_{c,i-1}, q'_{c,i}]$, therefore the probability for a coordinated user to have the highest U-value is $\frac{j}{j+s}$ where j is the number of the coordinated users with channel rate r_i. The proof for $P[C_{WIN}|MAX_{j,1}]$ is similar. □

Theorem 7.

$$P[C_{WIN}] = \sum_{i=1}^{M} \sum_{j=1}^{L} P[MAX_{j,i}] \cdot P[C_{WIN}|MAX_{j,i}],$$

where $P[MAX_{j,i}]$ and $P[COR|MAX_{j,i}]$ are given in Lemmas 2 and 3.

Proof: Immediate result from Bayes rule and the correctness of Lemmas 2 and 3. □

Theorem 8. *In a network consisting $L+N$ users (in the discrete model) of which L are coordinated, the average throughput (per time slot) of a coordinated user c is given by:*

$$\sum_{i=0}^{N} \binom{N}{i} \cdot \frac{1}{L} \sum_{i=1}^{M} \sum_{j=1}^{L} P[MAX_{j,i}] \cdot P[C_{WIN}|MAX_{j,i}] \cdot r_i \qquad (7)$$

where $P[MAX_{j,i}]$ and $P[COR|MAX_{j,i}]$ are given in Lemmas 2 and 3.

Proof: Let c be a coordinated user, the probability that some coordinated user obtains r_i at some time slot is given by $\sum_{j=1}^{L} P[MAX_{j,i}] \cdot P[C_{WIN}|MAX_{j,i}]$. The probability this user was c is $\frac{1}{L}$ and we get the probability he gets r_i is given by $\sum_{j=1}^{L} P[MAX_{j,i}] \cdot P[C_{WIN}|MAX_{j,i}] \cdot \frac{1}{L}$ and when summing it up over all possible rates we get Equation 7. □

4.5. Evaluation and Discussion

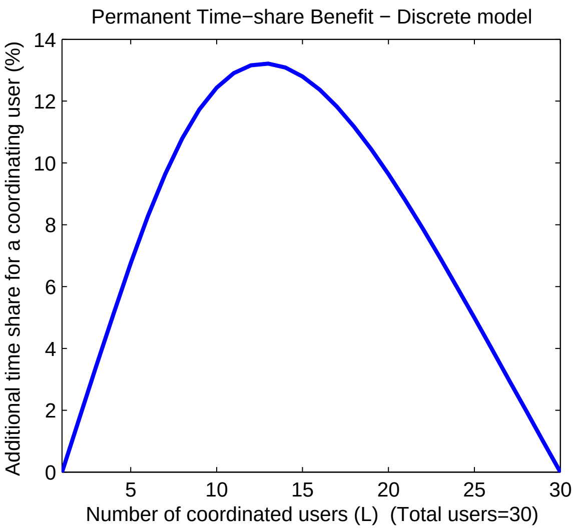

Figure 4: Y-axis is the additional time share (in percents) that a coordinated user (experiencing Rayleigh fading channel in CDMA2000 1xEV-DO) gains when he takes part in a coordinated group of size L (X-Axis). The results show that when there are 30 users in the system, participating in a coordinated group of 13 users is the most beneficial and increases the time share of the user by 13%.

While the share of a coordinated user in the continuous model depends only on the size of the coordinated group (L) and the number of regular users (N), the users' share in the discrete model depends also on the number of the possible channel rates (M) and their probabilities among the coordinated users $\{p_{c,i}\}_{i=1}^{M}$ (while still independent from the channel rate distributions of the regular users). The values of $\{p_{c,i}\}_{i=1}^{M}$ in the system configuration considered in figure 4 were set according to Rayleigh distribution on the 11 channel rates of the CDMA2000 1xEV-DO system in the same way which is already described in [17] (where the CDF scheduler for discrete channel rates was presented).

As in Figure 2, Figure 4 shows that the scheduler's notion of fairness is violated, we can see how the additional time share (Y-axis) for one coordinated user changes according to the number of users in the coordinated group (L) in the same manner as for the continuous model (Figure 2) and for the same reasons that were already mentioned in section 3.5. A remarkable difference between the models is that the time share benefit in the continuous model is greater than the benefit in the discrete model. By investigating the nature of the strategy effects in each model, we will be able to understand the reason for the differences and the effects of the strategy in different system configurations.

In both continuous and discrete models, when a coordinated user reports the minimal channel rate instead of his real channel rate then he widens the gap between the CDF of the distributions of his *real* (F_c) and *reported* (F'_c) channel rate distributions ($F'_c(r) > F_c(r)$). This gap defines the increase in the priority values the user gets for different rates. Therefore, an important observation is that the more the user gets to report a fake channel rates (according to the strategy conditions) - the more time slots he gets. In the continuous model, the shared information between the coordinated users allows them to identify in each time slot exactly $L-1$ users who have no chance winning while in the discrete model their number varies from 0 to $L-1$ (depending of how many coordinated users obtained the maximal channel rate $r_{m_c^*(t)}$). Therefore, there will always be more fake channel rate reports in the continuous model and the time share benefit in the continuous model will always be better than in the discrete model. As the number of the possible channel rates (M) grows, the coordinated users' behavior in the discrete model will become more like their behavior in the continuous model where the probability that two different users will get the same rate is zero. When M grows the probability that two users will get the same channel rate decreases. Therefore, less users obtain $r_{m_c^*(t)}$ and more users can report the minimal channel rate. This means that the benefit from the strategy will grow. Bigger sets of channel rates are expected in future physical standards for wireless communication to allow better utilization of channel fluctuations and/or to cover bigger range of channel condition values, so while it is expected to allow better system performance, it will make the system more vulnerable to such coordination strategy as we showed here.

Figure 5 shows the throughput benefit for the same settings as in Figure 4. In order to compute the throughput, channel rate probabilities $\{p_{c,i}\}_{i=1}^{M}$ were associated with actual rates in the CDMA2000 1xEV-DO as in [17]. For different sets of rates, we will get different results in Figure 5, while Figure 4 stays the same since it depends only on the set of probabilities $\{p_{c,i}\}_{i=1}^{M}$ regardless of the actual rates associated with them.

The throughput of a coordinated user was computed according to Theorem 8 and it was compared to the

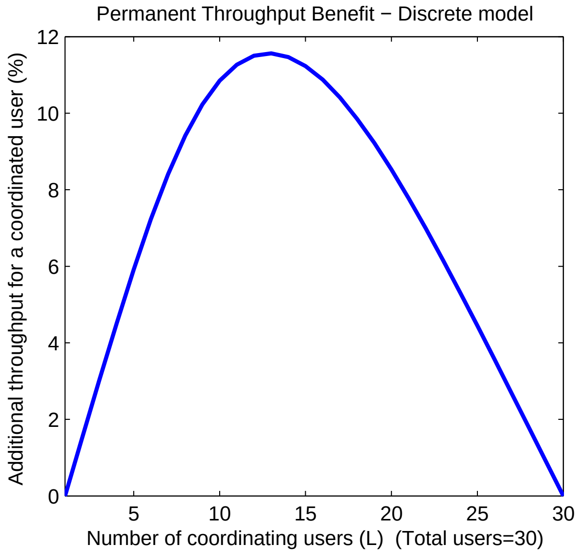

Figure 5: Y-axis is the additional throughput (in percents) that a coordinated user (in the same settings as in Figure 4) gains when he takes part in a coordinated group of size L (X-Axis). Like for time share, the results show that participating in a coordinated group of 13 users is the most beneficial, it increases the throughput of each user by 11.6%.

throughput that he gets under normal behavior which is given in [17]. Unsurprisingly, the throughput benefit demonstrates similar behavior to (as a function of L) the time share benefit.

5. System Loss by Malicious Strategy (Discrete Model)

In the previous section we focused on non-conformist opportunistic coordinated users whose objective is to increase their own share of the network resources. This increase was, of course, accompanied by performance degradation to the regular innocent users. This degradation can be easily computed from our results in Sections 3.5 and 4.5.

Our interest in this section is in malicious users whose objective is only to damage the other regular users, disregarding their own performance. So the malicious users are willing to degrade their own performance if it helps degrading that of the innocent users. Of course – the malicious users can damage the innocent users at least as much as opportunistic coordinated users can do. So the major question addressed in this section is whether they can inflict greater damage and how much.

Our focus in this analysis will be on the discrete distribution model which, due to its practicality, is of higher interest (than the continuous model) to system designers.

In section 5.1 we present a new strategy, the *malicious strategy* which allows malicious users to cause time share loss to the innocent users significantly higher than the loss caused by the *coordination strategy* (in the discrete model, Section 4). In fact, we prove that the damage caused by the malicious strategy under the discrete model is identical to the damage caused by the coordination strategy under the continuous model.

5.1. The Malicious Strategy

Malicious users intend to harm rather than increase their own throughput as other users in a coordinated group do. Hence, as we explain now, a malicious group has two advantages over a coordinated group of selfish users: 1. They gain a larger time share; 2. They need only to synchronize before starting the attack rather than exchanging information before every time slot as a selfish coordinated group does.

A coordinated user, aiming at increasing his own throughput, never reports a channel condition better than what he really experiences (if he does, it will decrease his expected throughput). Therefore, there may be some time slots where all the users in a coordinated group have exceptionally poor channel conditions and even the user (among them) with the best chances to be scheduled with the next time slot - has a very slim chance to get it. However, a malicious user, for whom his expected throughput is irrelevant, can report a channel condition independent of his real channel rate. This allows a malicious group to present in every time slot a user with an exceptionally good channel condition. For the same reason (irrelevancy of their throughput), the malicious users do not need to share any information on their real channel condition before every time slot and hence synchronizing once at the beginning of the attack is suffice, as demonstrated in the following description of the malicious strategy.

Assume a group of malicious users $MAL = \{1, 2, .., S\}$ (for simplicity we assume their indices are $1, 2, ..., S$). The basic idea followed by the malicious users is to take turns, in a round-robin fashion, in trying to obtain time slots. This means that at TS t, malicious user number (t mod S) will attempt to obtain the time slot. Consider a malicious user s: In all times slots that he tries to obtain he always reports the same channel rate - which we denote by r_{h_s} ($h \geq 2$), while in all other slots he reports other rates which are all lower than r_{h_s}. Each user s chooses his r_{h_s} value independently of other users. This flexibility in choosing channel rates makes the malicious pattern very hard to detect.

5.2. Analysis of the Malicious Strategy

Theorem 9. *In a network consisting $S + N$ users of which S are malicious which use the malicious strategy, the time share fraction dedicated to the S malicious users (jointly) depends only on S and N and equals*

$$\frac{S}{N+1}\left(1-\left(\frac{S-1}{S}\right)^{N+1}\right). \tag{8}$$

Proof: Consider a malicious user s with chosen maximal channel rate r_{h_s}. Since r_{h_s} is the highest channel rate he reports, then $q_{s,h_s} = 1$. Since in all other time slots - which are $\frac{S-1}{S}$ of the time - he reports r_{h_s-1} or lower, then we get $q_{s,h_s-1} = \frac{S-1}{S}$. Therefore, in every time slot t there will be *exactly* one malicious user, say s, with priority value $U_s \sim Uniform(\frac{S-1}{S}, 1)$.

Assume there are N regular users in the network. As stated in [17], the discrete scheduler preserves the fundamental character of the continuous model $U_n \sim Uniform(0,1)$ for every (regular) user n.

Let Z_i be the event where exactly i users get priority values in the range$(\frac{S-1}{S}, 1)$. Given Z_i, the probability that the malicious user will be the one to get the highest value in this range is $\frac{1}{i+1}$ since, given Z_i, the priority values of the malicious user and the regular users are uniformly distributed in $(\frac{S-1}{S}, 1)$. Unconditional on Z_i we get that the probability that a time slot is obtained by a malicious user is given by

$$\sum_{i=0}^{N} P[Z_i]\cdot\frac{1}{i+1}. \tag{9}$$

Now, $P[Z_i]$ is the probability that exactly $N-i$ users get priority value less than $\frac{S-1}{S}$, so according to $P[U_n < u] = u$ we get $P[Z_i] = \binom{N}{i}(\frac{S-1}{S})^{N-i}(\frac{1}{S})^i$. Thus, the probability that a malicious user wins a TS is

$$\sum_{i=0}^{N}\binom{N}{i}\left(\frac{S-1}{S}\right)^{N-i}\left(\frac{1}{S}\right)^{i}\cdot\frac{1}{i+1}. \tag{10}$$

Noting that Equation 10 is identical to Equation 5 (substitute S for L) we can use the analysis of Lemma 3 to obtain Equation 8. $\square$

Remark 1. *It is easy to see that the damage inflicted by this strategy on the innocent users (under the discrete model) is identical to the damage inflicted by coordinated users strategy under the continuous model.*

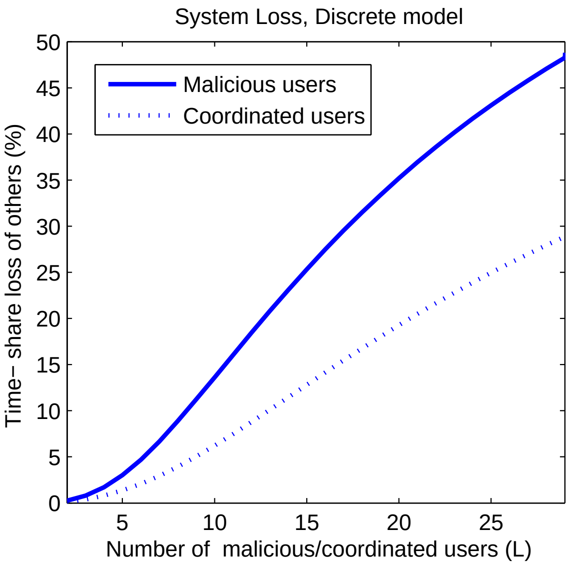

Figure 6: Time Slot loss of innocent users: Malicious vs. Coordination strategy. The network's population is 30 users.

5.3. Evaluation and Discussion

The effect of the malicious strategy is depicted in Figure 6. We consider a system consisting of 30 users in total and evaluate the (relative, percent-wise) time share loss experienced by each of the innocent users (compared to what he would get – 1/30 – in a normal system) as a function of the number of malicious users (x-Axis). This is done for both the coordinated user strategy (dotted line) and the malicious strategy (solid line). Both are evaluated for the discrete model. The channel distribution used is the Rayleigh distribution in CDMA2000 as in Figure 4. As one can observe, the loss caused by malicious users is significantly higher (by approximately a factor of 2) than that inflicted by the coordinated user strategy (for the discrete model). Note that according to Remark 1 we can use Corollary 4 to get an estimation for the time share loss of regular users. When $S = N - 1$ the time share of the regular users converges to e^{-1} instead of 0.5 means the loss of the regular users converges to 26% which is not far from the loss when $S = N = 15$ which equals 25% according to Figure 6.

Note that all the results in our work regarding malicious and coordinated users are independent of the channel rate distributions of regular users. While the throughput loss of of a regular user depends on his distribution, the case where some regular user experiences constant channel rate shows that the throughput loss can be high as the time share loss.

6. Solution Outline

The CDF scheduler is a unique scheduler. Almost magically, it maintains time share fairness without keeping track of past scheduling decisions. Unfortunately, while elegant, the CDF scheduler algorithm is too fragile. As we show in this work, the CDF scheduler fails to maintain fairness in the presence of selfish or malicious users. Hence, it has to be modified to take past scheduling decisions into account when making scheduling decisions. Such scheduler should record, for every user, the expected time share he should have got since he joined the system. For example, if a user has spent so far 200 time slots in the system with $N = 10$ users, it is expected that he should have received $1/N = 10\%$ of the time slots so far (20 time slots). The proportion between the expected time share so far of user k, denoted with $S^e_k(t)$, and the time share a user received in reality so far, denoted with $S^r_k(t)$, can be used to construct a weight that will influence his priority value positively or negatively - depends if the user is below or above his expected share. A simple example for such a modified scheduler can be as follows: Every user is assigned with a priority value as usual. Then, the priority value of every user k is multiplied with $w_k(t) = (S^e_k(t)/S^r_k(t))^\alpha$, where $\alpha > 1$ is constant decided by the system designer. It defines the balance between strict fairness enforcement and the overall throughput of the system. For example, when α is extremely high, then $w_k(t)$ is very close to zero if the user received more than he deserved $S^r_k(t) > S^e_k(t)$. In the same manner, $w_k(t)$ is very high for users who received less than they were expected ($S^r_k(t) < S^e_k(t)$). Observe that as the number of users in the system (N) increases, the expected fluctuation range of $S^e_k(t)/S^r_k(t)$ increases. Hence, a further refinement of such modified scheduler would be to use different α values for different values of N. That is, replace α with $\alpha(N)$. Note that the above is a very rough solution mainly meant to outline a possible solution. In our future work we aim at designing an optimal modified CDF scheduler which is immune to such attacks while maintaining an overall throughput in the system as close as possible to the original vulnerable CDF scheduler.

7. Conclusion

In this paper, based on scheme which targets the very fundamental principle of the CDF scheduler, we showed that non-conforming opportunistic users have the motivation to misreport their channel rates and destabilize the scheduler's notion of fairness. In addition we studied the loss for regular users inflicted by malicious users focused on degrading the system performance. We showed that for large populations consisting of regular and coordinated users in equal numbers, the ratio of allocated time slots between a coordinated and a regular user converges to $e - 1 \approx 1.7$. After researches proved the vulnerability of the Proportional Fairness scheduler, our work demonstrates the vulnerability of its alternative – the CDF scheduler. We recommend that this vulnerability, together with the solution we outlined should be taken into consideration by system designers when choosing and deploying a scheduler for modern wireless networks.

Appendix A. Misreporting of channel rates cannot benefit a single user

As promised, we provide a formal proof for this claim under the discrete model. Note that a similar proof can be constructed under the continuous model. Such proof will mainly differ in the throughput expressions and will use rates-ranges where single rates are used in the proof under the discrete model.

Lemma 4. *Let D_1 and D_2 be the following expressions*

$$\begin{aligned} D_1 &= \frac{(a-b)X + (b-c)Y}{a-c}(a^N - c^N) \\ D_2 &= X(a^N - b^N) + Y(b^N - c^N) \end{aligned}$$

where $a > b > c \geq 0$ and $N \geq 1$ is a natural number. Then, there exists a constant $d > 0$ such that $D_1 - D_2 = d(X - Y)$.

Lemma 4 can be easily proved using the identity $A^N - B^N = (A - B)\sum_{i=1}^{N} A^{N-i}B^{i-1}$. The full proof is provided in a technical report [18].

Theorem 10. *Under the discrete model, a user with no knowledge of the rates of other users cannot benefit from reporting fake channel rates.*

Proof: Both under the discrete and the continuous model, the priority values of all users are distributed uniformly in $[0, 1]$. Therefore, Theorem 1 is valid also under the discrete model. In order to complete the proof of the claim, we now prove that a user also cannot gain additional throughput by following a misreporting strategy. A user following a false-reports strategy sometimes report a certain channel rate which is different than his real channel rate. Let $p_{i,j}$ be the probability that in a random time slot the user reports r_i although his real rate is r_j. That is, $\sum_{i=1}^{M}\sum_{j=1}^{M} p_{i,j} = 1$. In order to evaluate the throughput a user gains when following a false-reports strategy, we first examine the outcome of winning a time slot with a fake report. When the user wins a time slot for which he reported to have a channel rate of r_i, the system tries to send him data at the rate he reported. If $r_i > r_j$, it means the real channel condition of the user can support data transfer only up to r_j. We assume, in favor[2] of false-reports strategies, that in such case the user receives r_j of the r_i the system sends him. If $r_i \leq r_j$, that is, the real channel rate of the user is good enough for the rate in which the system transmits (r_i). Formally, the expected rate received by a user if he wins a time slot for which he reported a channel condition of r_i while his real channel condition is r_j is given by $h_{i,j} = \min\{r_j, r_i\}$. Note that since we always discuss the same user k, for the sake of notational simplicity when we use q_i instead of $q_{k,i}$ until the end of the proof. The probability that a user reports r_i is $q_i - q_{i-1} = \sum_{j=1}^{M} p_{i,j}$. Therefore, given that the user won a time slot at which he reported r_i, the expected rate he receives is

$$H_i = \sum_{j=1}^{M} h_{i,j} p_{i,j} / (q_i - q_{i-1}). \tag{A.1}$$

Note that Eq. A.1 is valid only if $q_i - q_{i-1} > 0$. If $q_i - q_{i-1} = 0$ it means that the user never reports rate r_i. Hence, for the sake of completeness, we trivially define that $H_i = 0$ in this case. As explained earlier, in discrete CDF scheduling, when the user reports r_i he is assigned with a random number in $[q_{i-1}, q_i]$. In [17] the authors proved that $(q_i^N - q_{i-1}^N)/N$ (where N is the total number of users in the system) is the probability for a random slot to be: 1. A slot in which the user reports r_i and 2. a slot at which the user was assigned for transmission. Therefore, the expected throughput of a user is given by

$$T = \sum_{i=1}^{M} H_i (q_i^N - q_{i-1}^N)/N. \tag{A.2}$$

Observe that $H_i = r_i$ for a user who always report his real channel condition.

We now define a *Deviation Measurement* (DM) which evaluates (for a given strategy and user) how far from the truth are the reports of the user when he follows the strategy. The measure is given as follows: $DM(strategy, user) = \sum_{i=1}^{M} \sum_{j=1}^{M} p_{i,j} |i - j|$. Observe that if the user never fake his reports then $DM = 0$.

Assume by contradiction there exists some false-reports strategy that allows some user k (with some channel rate distribution) to achieve an expected throughput T_{high} higher than what he would achieve by always reporting his real rate. Let strategy $ST_{fake} \in \min_{st}\{DM(st, k) | T_{st} = T_{high}\}$ be a strategy with the minimal deviation among the strategies that achieve throughput T_{high}. We now prove that if ST_{fake} involves reporting a fake channel condition ($DM(fake, k) > 0$), there exists an alternative strategy ST_{alt} such that 1. $T_{alt} = T_{high}$; 2. $DM(ST_{alt}, k) < DM(ST_{fake}, k)$. This contradicts the definition of ST_{fake} and the assumption that throughput of T_{high} cannot be achieved by simply reporting the real channel condition. Let $\{\bar{p}_{i,j} | i, j \in [1, M]\}$ be the reporting probabilities of ST_{fake}. Since $DM(ST_{fake}, k) > 0$, at least one of the following claims has to be true: 1. There are $i > j$ such that $\bar{p}_{i,j} > 0$; 2. there are $i < j$ such that $\bar{p}_{i,j} > 0$. We first prove the existence of ST_{alt}, as described above, assuming claim 1. Let $w = \min_i\{\exists j < i . \bar{p}_{i,j} > 0\}$ and let $z = \min_j\{\bar{p}_{w,j} > 0\}$ (observe that $j < w$). Define an alternative strategy ST_{alt} as follows: ST_{alt} is identical to ST_{fake} with only one difference: At time slots in which ST_{fake} would instruct the user to report r_w instead of r_z, ST_{alt} instructs the user to report r_{w-1} instead. Note that by its definition, $w \geq 2$, hence r_{w-1} exists. Let $\{\tilde{p}_{i,j} | i, j \in [1, M]\}$ be the reporting probabilities when the user follows ST_{alt}. Then,

$$\tilde{p}_{i,j} = \begin{cases} i = w, j = z & 0 \\ i = w-1, j = z, & \bar{p}_{w-1,z} + \bar{p}_{w,z} \\ o/w & \bar{p}_{i,j} \end{cases} \tag{A.3}$$

Therefore, $DM(ST_{fake}, k) > DM(ST_{alt}, k)$. What is left to prove is that $T_{alt} \geq T_{fake}$. Let $\bar{H}_i$ ($\tilde{H}_i$) and $\bar{q}_i$ ($\tilde{q}_i$) be the expected received rates and the CDF values of ST_{fake} (ST_{alt}), Respectively. In S_{alt} the user sometimes reports r_{w-1} instead of r_w. Therefore:

$$\tilde{q}_i = \begin{cases} i = w-1 & \bar{q}_{w-1} + \bar{p}_{w,z} \\ o/w & \bar{q}_w \end{cases} \tag{A.4}$$

We now want to prove that $T^{alt} - T^{fake} \geq 0$. From Eq. A.2 we get that

$$T^{alt} - T^{fake} = \sum_{i=1}^{M} [\tilde{H}_i(\tilde{q}_i^N - \tilde{q}_{i-1}^N) - \bar{H}_i(\bar{q}_i^N - \bar{q}_{i-1}^N)] \tag{A.5}$$

Both strategies are similar when it comes to reports of rates different than r_w and r_{w-1}. Therefore, $\forall i \neq w, w-1 . \tilde{H}_i = \bar{H}_i$ (Formally, that can be concluded from equations A.1, A.3 and A.4). Considering also Eq. A.4,

[2]The user might receive less than r_j since the transmission modulation might not be the optimal one for his real channel condition.

we can contract Eq. A.5 to:

$$T^{alt} - T^{fake} = \qquad (A.6)$$
$$\tilde{H}_w(\tilde{q}_w^N - \tilde{q}_{w-1}^N) - \bar{H}_w(\bar{q}_w^N - \bar{q}_{w-1}^N) +$$
$$\tilde{H}_{w-1}(\tilde{q}_{w-1}^N - \tilde{q}_{w-2}^N) - \bar{H}_{w-1}(\bar{q}_{w-1}^N - \bar{q}_{w-2}^N)$$

Using equations A.3, A.4 and A.1 Eq. A.6 can be writtenas follows[3]:

$$\begin{aligned} T^{alt} - T^{fake} &= \tilde{H}_w(\bar{q}_w^N - (\bar{q}_{w-1} + \bar{p}_{w,z})^N) \\ &- \frac{\tilde{H}_w(\bar{q}_w - (\bar{q}_{w-1} + \bar{p}_{w,z})) + r_z\bar{p}_{w,z}}{\bar{q}_w - \bar{q}_{w-1}}(\bar{q}_w^N - \bar{q}_{w-1}^N) \quad (A.7) \\ &+ \frac{\bar{H}_{w-1}(\bar{q}_{w-1} - \bar{q}_{w-2}) + r_z\bar{p}_{w,z}}{\bar{q}_{w-1} + \bar{p}_{w,z} - \bar{q}_{w-2}}((\bar{q}_{w-1} + \bar{p}_{w,z})^N - \bar{q}_{w-2}^N) \\ &- \bar{H}_{w-1}(\bar{q}_{w-1}^N - \bar{q}_{w-2}^N) \end{aligned}$$

Note that since $\bar{p}_{w,z} > 0$ the denominators in Eq. A.7 must be greater than zero. In addition, note that is it possible that both or one of $\tilde{H}_w$ and $\bar{H}_{w-1}$ is zero. We now examine Eq. A.7 and prove that:

$$\begin{aligned} &\frac{\bar{H}_{w-1}(\bar{q}_{w-1} - \bar{q}_{w-2}) + r_z\bar{p}_{w,z}}{\bar{q}_{w-1} + \bar{p}_{w,z} - \bar{q}_{w-2}}((\bar{q}_{w-1} + \bar{p}_{w,z})^N - (\bar{q}_{w-2})^N) \\ &- \bar{H}_{w-1}((\bar{q}_{w-1})^N - (\bar{q}_{w-2})^N) \quad (A.8) \\ &\geq r_z((\bar{q}_{w-1} + \bar{p}_{w,z})^N - (\bar{q}_{w-1})^N) \end{aligned}$$

If $\bar{H}_{w-1} = 0$ (and therefore $\bar{q}_{w-1} = \bar{q}_{w-2}$) then Eq. A.8 is true. If $\bar{H}_{w-1} > 0$, then Eq. A.8 is a result of Lemma 4 as follows. Let $a = \bar{q}_{w-1} + \bar{p}_{w,z}$, $b = \bar{q}_{w-1}$ and $c = \bar{q}_{w-2}$. Note that $\bar{p}_{w,z} > 0$ hence $a > b$. In addition, $b > c$ since $\bar{H}_{w-1} > 0$. Define $X = r_z$ and $Y = \bar{H}_{w-1}$. Recall that r_w is the lowest channel rate the user reports (in ST_{fake}) while his real channel rate is lower. Therefore, when the user reports r_{w-1} in ST_{fake}, his real rate has to be equal or higher than r_{w-1}. Therefore, $\forall j.\ \ \bar{p}_{w-1,j} \neq 0 \rightarrow h_{w-1,j} = r_{w-1}$ and we get that $Y = \bar{H}_{w-1} = r_{w-1}$. In addition, recall that $z < w$ and hence, $Y = r_{w-1} \geq r_z = X$.

Therefore, in the context of Lemma 4 we get that $D_1 \geq D_2$ and therefore $\frac{(a-b)X+(b-c)Y}{a-c}(a^N - c^N) - Y(b^N - c^N) \geq X(a^N - b^N)$, which is identical to Eq. A.8 if replacing a, b, c, X and Y with the values with which they were defined.

From equations A.8 and A.7 we get that

$$\begin{aligned} T^{alt} - T^{fake} &\geq \qquad (A.9) \\ &\tilde{H}_w((q_w)^N - (q_{w-1} + p_{w,z})^N) \\ &- \frac{\tilde{H}_w(q_w - (q_{w-1} + p_{w,z})) + r_z p_{w,z}}{q_w - q_{w-1}}((q_w)^N - (q_{w-1})^N) \\ &+ r_z((q_{w-1} + p_{w,z})^N - (q_{w-1})^N) \end{aligned}$$

[3] A detailed explanation can be found in the technical report [18].

If $\tilde{H}_w = 0$ we immediately get that $T^{alt} - T^{fake} \geq 0$. Otherwise, we use Lemma 4 to prove this claim. This time, in the context of Lemma 4, we define $a = q_w$, $b = q_{w-1} + p_{w,z}$, $c = q_{w-1}$, $X = \tilde{H}_w$ and $Y = r_z$. Note that $b > c$ since $p_{w,z} > 0$ and $a > b$ since $\tilde{H}_w > 0$.

Observe that in the context of Lemma 4 we can write equation A.9as follows:

$$T^{alt} - T^{fake} \geq D_2 - D_1 \qquad (A.10)$$

Recall the way z was chosen. In strategy ST_{fake}, r_z is the lowest possible rate the user experiences if he reports r_w. Therefore, in strategy ST_{alt}, the lowest real rate of the user when he reports r_w cannot be lower than r_{z+1}. Therefore, $X = \tilde{H}_w > r_z = Y$ and we get that $D_2 - D_1 > 0$ and hence proved that $T^{alt} \geq T^{fake}$.

Recall that we explained that since $DM(ST_{fake}, k) > 0$, at least one of the following claims has to be true: 1. There are $i > j$ such that $\bar{p}_{i,j} > 0$; 2. there are $i < j$ such that $\bar{p}_{i,j} > 0$. We have just proved that if claim 1 is true then there is a strategy ST_{alt} such that $T^{alt} \geq T^{fake}$ and $DM(ST_{fake}, k) > DM(ST_{alt}, k)$ which contradicts the definition of ST_{fake}. Now all is left to prove is that if claim 2 is true, there exists strategy ST_{alt2} such that $T^{alt2} \geq T^{fake}$ and $DM(ST_{fake}, k) > DM(ST_{alt2}, k)$. Let $w' = \max_i\{\exists j > i.\bar{p}_{i,j} > 0\}$ and let $z' = \max_j\{\bar{p}_{w',j} > 0\}$ (observe that $j > w'$). Define an alternative strategy ST_{alt2} as follows: ST_{alt2} is identical to ST_{fake} with only one difference: At time slots in which ST_{fake} would instruct the user to report $r_{w'}$ instead of $r_{z'j}$, ST_{alt2} instructs the user to report $r_{w'+1}$ instead. (Note that by its definition, $w' \leq M - 1$, hence $r_{w'+1}$ exists.). The strategy ST_{alt2} and the proofs for the above claims are completely symmetrical to ST_{alt} and the proofs of its properties. Hence, due to luck of space we exclude from this article the complete proof. It can be found in a technical report [18]. □

References

[1] H. Holma, A. Toskala, HSDPA/HSUPA for UMTS: High Speed Radio Access for Mobile Communications, John Wiley & Sons, 2006.

[2] T. I. Association, CDMA2000: High Rate Packet Data Air Interface Specification (TIA-856-A).

[3] V. Vanghi, A. Damnjanovic, B. Vojcic, The CDMA2000 System for Mobile Communications: 3G Wireless Evolution, Prentice Hall PTR, Upper Saddle River, NJ, USA, 2004.

[4] R. P. A. Jalali, R. Pankaj, Data Throughput of CDMA-HDR: A High Efficiency High Data Rate Personal Communication Wireless System, in: Proceedings of the IEEE Vehicular Technology Conference.

[5] S. Borst, User-Level Performance of Channel-Aware Scheduling Algorithms in Wireless Data Networks, IEEE/ACM Trans. Netw. 13 (2005) 636–647.

[6] D. Park, H. Seo, H. Kwon, B. G. Lee, Wireless Packet Scheduling Based on the Cumulative Distribution Function of User Transmission Rates, IEEE Transactions on Communications 53 (2005) 1919–1929.

[7] S. Bali, S. Machiraju, H. Zang, V. Frost, A Measurement Study of Scheduler-Based Attacks in 3G Wireless Networks, in: PAM, pp. 105–114.

[8] R. Racic, D. Ma, H. Chen, X. Liu, Exploiting Opportunistic Scheduling in Cellular Data Networks, in: Proceedings of the 15th Annual Network and Distributed System Security Symposium (NDSS 2008), San Diego, CA.

[9] M. Guirguis, A. Bestavros, I. Matta, Exploiting the Transients of Adaptation for RoQ Attacks on Internet Resources, in: in Proceedings of the 12th IEEE International Conference on Network Protocols (ICNP04).

[10] M. Guirguis, A. Bestavros, I. Matta, Y. Zhang, Reduction of Quality (RoQ) Attacks on Internet End-Systems, in: in Proceedings of Infocom05: The IEEE International Conference on Computer Communication, pp. 1362–1372.

[11] R. Smith, C. Estan, S. Jha, Backtracking Algorithmic Complexity Attacks Against a NIDS, in: ACSAC.

[12] A. Kuzmanovic, E. W. Knightly, Low-Rate TCP-Targeted Denial of Service Attacks (The Shrew VS. the Mice and Elephants), in: ACM SIGCOMM Conference on Applications.

[13] A. Bremler-Barr, H. Levy, N. Halachmi, Aggressiveness Protective Fair Queueing for Bursty Applications, in: IWQoS.

[14] E. Doron, A. Wool, Wda: A Web Farm Distributed Denial of Service Attack Attenuator, Comput. Netw. 55 (2011) 1037–1051.

[15] C. Castelluccia, E. Mykletun, G. Tsudik, Improving Secure Server Performance by Re-balancing SSL/TLS Handshakes, in: USENIX.

[16] U. Ben-Porat, A. Bremler-Barr, H. Levy, On the Exploitation of CDF Based Wireless Scheduling, in: INFOCOM, pp. 2821–2825.

[17] D. Park, B. G. Lee, Qos Support by Using CDF-Based Wireless Packet Scheduling in Fading Channels, IEEE Transactions on Communications 54 (2006) 955.

[18] U. Ben-Porat, A. Bremler-Barr, H. Levy, On the Exploitation of CDF Based Wireless Scheduling, 2012. Http://www.faculty.idc.ac.il/bremler/.