The 18th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'07)

Opportunistic Scheduling with Statistical Fairness Guarantee in Wireless Networks

[1]Yu Wu, [1]Ji-bo Wei, [1]Yong Xi, [2]Byung-Seo Kim and [3]Sung Won Kim
[1]School of Electronic Science and Engineering, National University of Defense Technology, Changsha, Hunan, 410073 China
[2]Networks and Enterprise, Motorola Inc, Schaumburg, IL, 60196 USA
[3]School of Electrical Engineering and Computer Science, Yeungnam University, Gyeong-San, Gyeongbuk, 712749 Korea

ABSTRACT

In this paper, we present Statistical Fairness Opportunistic Scheduling (SFOS), a wireless scheduling algorithm with the objective of improving system throughput while providing statistical fairness guarantee. In particular, SFOS provides statistical fairness by using a virtual time variation, while improving system throughput by using our designed *utility function* which relates to the transmission rate. We develop a general analytical framework for SFOS, which shows that the fairness index is bounded by the utility. Further, we investigate the design rule of the utility function over Rayleigh fading channels and discrete transmission rates and present a reference design. Simulation results evaluate SFOS can significantly improve system throughput while providing statistical fairness guarantee.

I. INTRODUCTION

Forthcoming wireless systems such as the fourth generation (4G) cellular systems are supposed to support a variety of services, requiring different Quality of Service (QoS). Since these services may require significantly large amount of data to be delivered, along with the low reliability and time varying capacity of the wireless channels, efficient bandwidth allocation and scheduling scheme is a priority.

Opportunistic Scheduling (OS) is a modern approach of improving the performance of wireless networks. When making selection, OS selects the user which has favorable channel conditions to transmit. In this context, the more users to choose from and the more randomness the channel exhibits, the larger the gain the approach can get. Meanwhile, though a greedy OS scheme that always schedules the user with the highest data rate may maximize the channel utilization, it may lead to serve unfairness. Hence, the fairness issues have been considered in almost all of the existing OS schemes [1-3]. However, until recently, the "fairness criteria" defined by prior approaches are still rather arbitrary and seldom reflect the actual requirement of the system. For example, the fairness objective of the Proportional Fair (PF) scheduler [3] is to asymptotically provide equal service time across all users in the system, which may be too rigid fairness constraint for different services and prevent further improving system throughput. The Wireless Credit-based Fair Queuing (WCFQ) presented in [4] is a well-devised OS scheme that allows system operators have flexibility in tradeoff between perfect temporal fairness and purely opportunistic best-user scheduling. However, this algorithm is designed on the assumptions that the channel condition has a uniform distribution and a linear dependency on the transmission rate, which are far from the actual situations of current wireless networks. Therefore, the scheme is impractical in applications.

In this paper, we present Statistical Fairness Opportunistic Scheduling (SFOS), a wireless scheduler based on the statistical fairness principle proposed in [4]. In SFOS, the virtual time originally proposed in [5] is redefined to provide temporal fairness guarantee, while our defined *utility function* is devised to improve system throughput. Furthermore, we resolve many critical problems when applying SFOS over Rayleigh fading channels and finite transmission rates. Simulation results evaluate that SFOS can obviously improve system throughput while satisfying the practical fairness requirement of the system.

The rest of the paper is organized as follows. Section II describes the system model and the design objective of SFOS. Section III presents the SFOS algorithm and analyzes its fairness properties. Section IV presents the design method of the utility function. Section V shows the simulation results. Section VI concludes this paper.

II. SYSTEM MODEL AND DESIGN OBJECTIVE

A. System Model

We consider a time-slotted cellular system consisting of one BS and a number of flows. The system has only one channel, thus at any time only one flow can transmit in the cell. Scheduling is conducted by the BS. Here we only focus on the performance of downlink scheduling. We assume the timeslot width is fixed and the physical frame size can be varied according to the transmission rate so that each transmission occurs at the beginning of each timeslot as assumed in [6-8]. We further assume the channel conditions (represented by the received SNR) of uses are constant over a timeslot. Since our scheduler will use channel conditions to make decision, we assume some additional mechanisms can be utilized to make predicted channel conditions available to the BS.

B. Design Objective of SFOS

The design objective of SFOS is to improve system throughput while providing statistical temporal fairness guarantee. In this paper, we use the definition of system throughput proposed in [2] as follows:

$$S = \lim_{T\to\infty} \frac{1}{T} \sum_{i=1}^{N} \sum_{t=0}^{T-1} R_i(t) \cdot I_i(t), \qquad (1)$$

where N is the number of all the backlogged flows in the system at timeslot t. $R_i(t)$ is the transmission rate of flow i at timeslot t. $I_i(t) = 1$ if flow i is scheduled at timeslot t, otherwise,

$I_i(t)=0$. Since at any timeslot, only one flow can be scheduled, $\sum_{i=1}^{N} I_i(t)=1$. Now we formally present the fairness objective of SFOS as follows:

$$\Pr\left(\left| \frac{W_i(t_1,t_2)}{\phi_i} - \frac{W_j(t_1,t_2)}{\phi_j} \right| \geq x \right) \leq g(i,j,x), \quad (2)$$

where $W_i(t_1,t_2)$ represents the service time (in timeslot) received by flow i over any time interval (t_1,t_2), ϕ_i is flow i's weight assigned by the system, $g(i,j,x)$ is our defined *fairness function*, which decreases with the weighted service discrepancy, x, in timeslot. Therefore Ineq. (2) defines statistical fairness in terms of the distribution of the weighted service discrepancy between any two flows in the system. We present such a fairness guarantee based on two considerations. First, it is more practical to provide statistical fairness guarantee to users rather than a deterministic bound because of the randomness inherent in wireless channels. Second, since the fairness guarantee is now bounded by the fairness function, system operators can accurately reflect current fairness requirement of the system by designing this function.

III. The SFOS Algorithm

A. Scheduling Policy

We start by introducing the notations. Let $B(t)$ be the set of all the backlogged flows in the system at timeslot t, $U_i(t)$ be our defined *utility function*, which increases with the transmission rate of flow i at timeslot t. To maintain fairness, we adopt Start-time Fair Queueing (SFQ) [5]. Thus, SFOS maintains a system virtual time, $V(t)$, in timeslot. In addition, it associates each flow i with a virtual time $V_i(t)$. Intuitively, $V(t)$ represents the minimum service time each flow should have received by timeslot t, and $V_i(t)$ records the service time flow i has received by timeslot t. Table I shows the SFOS algorithm and the description is as follows.

Table I. Pseudo of the SFOS algorithm

```
enqueue flow i's packet:
1:  if (i ∉ B(t))
2:     V_i(t) = max(V(t), V_i(t-1));
dequeue():
3:  i = min{ V_k(t) + 1/φ_k − U_k(t) | k ∈ B(t), R_k(t) > 0 };
4:  if (i exists)
5:     V_i(t+1) = V_i(t) + 1/φ_i;
6:     for (k ∈ B(t)\{i})
7:        V_k(t+1) = V_k(t);
8:     V(t+1) = min{ V_k(t) | k ∈ B(t) };
```

When a flow i becomes backlogged, its virtual time is initialized to the maximum of current system virtual time and its previous virtual time (Line 2). This step can ensure the virtual times of all the backlogged flows are bounded. The most important procedures are scheduling decision. Other than selecting a flow with the minimum virtual time in order to obtain absolute fairness like SFQ, SFOS incorporates channel conditions into the scheduling decision. Therefore, both utility and virtual time directly affect the selection result (Line 3). Specifically, if a flow has much bigger utility, it can be scheduled even it has accumulated a large amount of virtual time. In this way, system throughput gets improved. On the other hand, an under-serviced flow can get compensated once its virtual time becomes relatively much smaller even at that time it has small utility.

Following scheduling decision, the virtual time of the selected flow is increased (Line 5). Finally, the system virtual time is updated to the minimum of the virtual times of all the backlogged flows in the system (Line 8).

B. Performance Analysis

We first introduce Lemma 1 that the difference between the virtual times of any two backlogged and error-free flows is bounded by the utility. We then derive their statistical bounds of the weighted service discrepancy in Theorem 1.

Lemma 1 The difference between the virtual times of any two backlogged and error-free flows i and j at timeslot t is bounded as follows:

$$-\frac{1}{\phi_i} + U_i(t) - U_j(t) \leq V_i(t+1) - V_j(t+1) \leq \frac{1}{\phi_j} + U_i(t) - U_j(t), \quad (3)$$

Proof: First, assume flow i is selected at timeslot t, then according to the selection rule of SFOS, we have:

$$V_i(t) + \frac{1}{\phi_i} - U_i(t) \leq V_j(t) + \frac{1}{\phi_j} - U_j(t). \quad (4)$$

From the update method of the virtual time, we have:

$$V_i(t+1) = V_i(t) + \frac{1}{\phi_i} \quad (5)$$
$$V_j(t+1) = V_j(t),$$

Bring these two updated virtual times into (4), we have:

$$V_i(t+1) - V_j(t+1) \leq \frac{1}{\phi_j} + U_i(t) - U_j(t). \quad (6)$$

Similarly, if assume flow j is selected at timeslot t, we have:

$$V_i(t+1) - V_j(t+1) \geq -\frac{1}{\phi_i} + U_i(t) - U_j(t). \quad (7)$$

Thus from (6) and (7), we conclude the proof. □

Theorem 1: For any two backlogged and error-free flows i and j over any time interval (t_1,t_2), if each flow's utility is independent identically distributed (i.i.d.), then we have the statistical fairness guarantee for their weighted service discrepancy as follows:

$$\Pr\left(\left|\frac{W_i(t_1,t_2)}{\phi_i}-\frac{W_j(t_1,t_2)}{\phi_j}\right|\geq\frac{2}{\phi_i}+\frac{2}{\phi_j}+4x\right)\leq\Pr(U_i+U_j\geq x+x). \quad (8)$$

Proof: Considering for any backlogged and error-free flow i, its received service time over any time interval (t_1,t_2) has the following bounds [9]:

$$V_i(t_2)-V_i(t_1)-\frac{1}{\phi_i}\leq\frac{W_i(t_1,t_2)}{\phi_i}\leq V_i(t_2)-V_i(t_1)+\frac{1}{\phi_i}. \quad (9)$$

Thus, the weighted service discrepancy can be represented by the virtual time as follows:

$$\begin{aligned}V_i(t_2)-V_j(t_2)-(V_i(t_1)-V_j(t_1))-\frac{1}{\phi_i}-\frac{1}{\phi_j}&\leq\frac{W_i(t_1,t_2)}{\phi_i}-\frac{W_j(t_1,t_2)}{\phi_j} \\ &\leq V_i(t_2)-V_j(t_2)-(V_i(t_1)-V_j(t_1))+\frac{1}{\phi_i}+\frac{1}{\phi_j},\end{aligned} \quad (10)$$

where, j is any other backlogged and error-free flow over (t_1,t_2). Then from **Lemma 1**, we have:

$$\begin{aligned}-\frac{1}{\phi_i}-\frac{1}{\phi_j}+\Delta_{i,j}(t_2)-\Delta_{i,j}(t_1)&\leq V_i(t_2)-V_j(t_2)-(V_i(t_1)-V_j(t_1)) \\ &\leq\frac{1}{\phi_i}+\frac{1}{\phi_j}+\Delta_{i,j}(t_2)-\Delta_{i,j}(t_1),\end{aligned} \quad (11)$$

where $\Delta_{i,j}(t_1)=U_i(t_1)-U_j(t_1)$. Bring (11) into (10) and let $U_i(\lambda)=\max\{|U_i(t_1)|,|U_i(t_2)|\}$, $U_j(\varsigma)=\max\{|U_j(t_1)|,|U_j(t_2)|\}$, it immediately follows:

$$\left|\frac{W_i(t_1,t_2)}{\phi_i}-\frac{W_j(t_1,t_2)}{\phi_j}\right|\leq\frac{2}{\phi_i}+2U_i(\lambda)+\frac{2}{\phi_j}+2U_j(\varsigma). \quad (12)$$

Further, by relaxing the discrepancy using (12), we have:

$$\Pr\left(\left|\frac{W_i(t_1,t_2)}{\phi_i}-\frac{W_j(t_1,t_2)}{\phi_j}\right|\geq\frac{2}{\phi_i}+\frac{2}{\phi_j}+4x\right)\leq\Pr(U_i(\lambda)+U_j(\varsigma)\geq 2x). \quad (13)$$

Thus from the i.i.d. assumption, Theorem 1 holds. □

Theorem 1 shows the fairness guarantee of SFOS is related to the utility. When the utility is small, the fairness constraint is correspondingly tight, and the weighted service discrepancy of any two flows is very small. As the utility increases, the fairness constraint becomes weaker and the utility will dominate the scheduling selection. As a result, the flows with larger utility will receive more transmissions. Since the utility of each flow is related to its transmission rate, it is easy to see that system throughput will increase with the increasing utility.

IV. Design of The Utility Function

In this section, we present the design methodology of the utility function. But we first discuss the design rule of the fairness function of the system.

A. Design of Fairness Function

We assume the system has the following statistical short-term temporal fairness requirement:

$$\Pr\left(\left|\frac{W_i(t_1,t_2)}{\phi_i}-\frac{W_j(t_1,t_2)}{\phi_j}\right|\geq\frac{2}{\phi_i}+\frac{2}{\phi_j}+4x\right)\leq g(x). \quad (14)$$

where, $g(x)$ is the fairness function of the system. Now observe (14), $g(x)$ must have the following basic properties. First, $g(x)$ should be within [0,1] for $x\geq-\frac{1}{2}(1/\phi_i+1/\phi_j)\triangleq-L$. Second, $g(x)$ should decrease towards 0 with increasing x in order to make the statistical fairness constraint much tight. Hence, according to Theorem 1, to satisfy Ineq. (14), as long as the utility function satisfies:

$$\Pr(U_i+U_j\geq x+x)\leq g(x). \quad (15)$$

Considering:

$$\Pr(U_i+U_j\leq x+x)\geq\Pr(U_i\leq x)\cdot\Pr(U_j\leq x),$$

Ineq. (15) can be satisfied as long as:

$$\Pr(U_i\leq x)\cdot\Pr(U_j\leq x)\geq 1-g(x). \quad (16)$$

To further simplify designing, we present a reference design of $g(x)$ as follows:

$$\begin{aligned}g(x)&=1-(1-\exp(-\frac{x+L}{\beta}\cdot\frac{P_0}{\overline{P_i}}))\cdot(1-\exp(-\frac{x+L}{\beta}\cdot\frac{P_0}{\overline{P_j}})) \\ &\triangleq 1-\left(1-g_i(x)\right)\cdot\left(1-g_j(x)\right), \qquad \beta,P_0\geq 0\end{aligned} \quad (17)$$

where, $\overline{P_i}$ is the average received power (in W) of flow i. P_0 is a constant to normalize $\overline{P_i}$. β is a tunable parameter. Then from (16), we have:

$$\begin{aligned}\Pr(U_i\leq x)&\geq 1-g_i(x) \\ \Pr(U_j\leq x)&\geq 1-g_j(x).\end{aligned} \quad (18)$$

Compared to (16), Ineqs. (18) share the fairness constraint of the system on each flow, thus we can design each flow's utility function according to its assigned fairness constraint.

B. Design of Utility Function

From Ineqs. (18) we find that the expression of the utility function can be derived if we find the distribution function of the utility. In the following, we discuss this method.

Here, we select BPSK, QPSK, QAM16, QAM64 and QAM256 as the modulation schemes in the system, but not use any correlation coding scheme. With channel bandwidth set to 1MHz, the highest transmission rates supported by these modulation schemes while satisfying the bit error rate (BER) to 1e-5 are listed as follows:

$$\begin{cases}\text{BPSK (1Mbps)}: & SNR\geq 9.6\text{ dB} \\ \text{QPSK (2Mbps)}: & SNR\geq 12.6\text{ dB} \\ \text{QAM16 (4Mbps)}: & SNR\geq 19.5\text{ dB} \\ \text{QAM64 (6Mbps)}: & SNR\geq 25.6\text{ dB} \\ \text{QAM256 (8Mbps)}: & SNR\geq 31.5\text{ dB}\end{cases} \quad (19)$$

Obviously, the transmission rate is a step function of the received SNR. Then based on (18) and (19), we can use the distribution function of the receive SNR to represent the distribution function of the utility as follows:

$$
\begin{aligned}
&\text{1) when } f^{-1}(x) < 1, 2, 4, 6, 8\text{Mbps respectively}\\
&\quad \Pr(U_i \le x) = \Pr(R_i \le f^{-1}(x)) = \Pr(\rho_i < \rho) \ge 1 - g_i(x)\\
&\quad \text{where, } \rho = 9.6, 12.6, 19.5, 25.6, 31.5 \text{ dB respectively}\\
&\text{2) when } f^{-1}(x) < \infty\\
&\quad \Pr(U_i \le x) = \Pr(R_i \le f^{-1}(x)) = 1 \ge 1 - g_i(x),
\end{aligned} \tag{20}
$$

where ρ_i is the received SNR (in dB) of flow i. Then if Rayleigh fading channels are assumed, we can derive the distribution function of ρ_i as follows:

$$\Pr(\rho_i < \rho) = 1 - \exp(-\frac{\Gamma_i}{\overline{P}_i} \cdot 10^{\rho/10}), \tag{21}$$

where Γ_i is the expected power of white Gaussian noise (in W) of flow i, $\overline{P}_i$=E(y_i) is the mean of the received power (in W) of flow i. Bring (21) into (20), we can derive the expression of the utility function as follows:

$$
\begin{aligned}
&\text{1) when } R_i = 0, 1, 2, 4, 6\text{Mbps respectively}\\
&\quad U_i = \frac{\beta \cdot \Gamma_i}{P_0} \cdot 10^{\rho/10} - L,\\
&\quad \text{where, } \rho = 9.6, 12.6, 19.5, 25.6, 31.5 \text{ dB respectively}\\
&\text{2) when } R_i = 8\text{Mbps}\\
&\quad U_i = \infty.
\end{aligned} \tag{22}
$$

In implementation, when R_i=8Mbps, set U_i to a value bigger than $\frac{\beta \cdot \Gamma_i}{P_0} \cdot 10^{31.5/10} - L$.

From (22) and (17), we make two observations. First, both utility function and fairness function are determined by β. This makes system operators can flexibly change the fairness constraint by setting different β values. For example, if the system requires strong fairness guarantee, β should be set to a small value which leads to a small utility function and thus the virtual time will dominate the scheduling selection. Otherwise, if the system allows weak guarantee, in other word, it favors higher system throughput, β can be set to a big value which makes the utility function dominate the scheduling selection. Second, when given β, the utility function of each flow only varies with its transmission rate but irrespective of its average received power. This is just the reason we actively introduce $\overline{P}_i$ into fairness function (see (17)). Otherwise, the average power will exist in the utility function since it directly influences the distribution function of the received SNR (see (21)). This will ultimately cause the scheduler biases towards the flows with smaller average power and tends to make system throughput decreased. Since this problem always exists over Rayleigh fading channels and discrete transmission rates, we may have to address it by carefully designing the fairness function.

V. Simulation Evaluations

In this section, we conduct an extensive suit of simulation to evaluate SFOS. In comparison, we also simulate MR-FQ [10], PF and greedy OS. Since MR-FQ provides identical service time to all the users in the system, it is used as fairness benchmark, whereas the greedy OS acts as system throughput benchmark. For ease of interpretation, we simulate only two users and each of them receives a continuously backlogged Constant Bit Rate (CBR) flow. The weight of each flow is 1. Both users keep static during the whole simulation and their distances from the BS are 50m and 150m respectively. The relationship between received SNR and highest transmission rate consists with (19). As a result, their average transmission rates are 8Mbps and 2Mbps respectively.

It is worth noticing that we should evaluate the fairness properties of SFOS in an error-free system according to Theorem 1. However, it is impractical for Rayleigh fading channels. To address this problem, we make modifications to the simulated channels to allow the selected user can still transmit with the smallest transmission rate even its current received SNR is below the lowest threshold.

A. System Throughput and Long-Term fairness

Fig. 1 plots system throughput under different schemes, while Fig. 2 plots the temporal share received by each user during the whole simulation run (10,000 timeslots), wherein the bottom part of the bars represents the temporal share of the user closer to the BS.

We can see from Fig. 1 that system throughput increases with β. Specially, when β=0, system throughput under SFOS is almost the same as that under MR-FQ. The reason is at this time both of them make scheduling only based on the virtual time but irrespective of the transmission rate. Thus from Fig. 2, we can see each user receives absolutely equal share. Thereafter, system throughput under SFOS gradually increases with β. This means the impact of the utility on the scheduling decision keeps increasing and the user with higher transmission rate gets more chance to transmit. For example, the temporal share received by the user closer to the BS is 1.47 times than the other when β=2^8, while 2.35 times when β=2^9. From β=2^2, system throughput under SFOS begins to exceed that under PF. This is because the fairness objective of PF is to provide equal temporal shares to all the users (see Fig. 2), which restrains system throughput from further improving especially when the users have distinct channel conditions. Fig. 1 shows, when β=2^{13} system throughput under SFOS almost approaches that under the greedy OS, which is 17.3% more than that under MR-FQ and 12.0% more than that under PF. Meanwhile, we can see from Fig. 2 that the user closer to the BS receives 96.7% of the total service time. This means at this time SFOS almost makes selection only based on transmission rate like greedy OS.

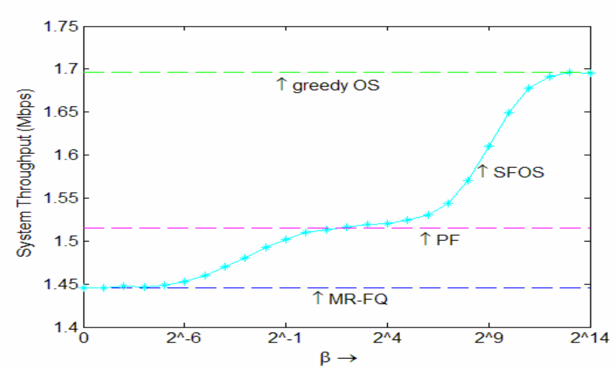

Fig. 1. System throughput under different schemes

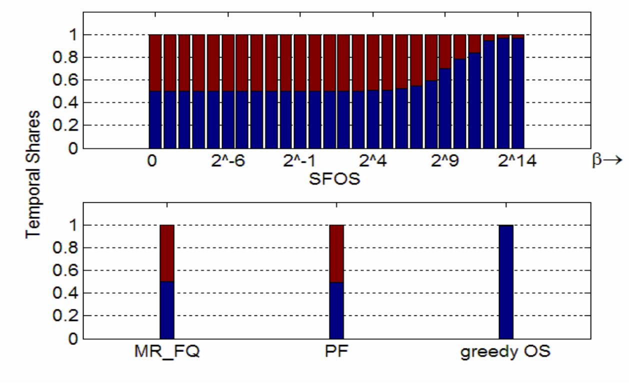

Fig. 2. Temporal shares over 10,000 timeslots

B. *Short-term Fairness Guarantee*

Here, we evaluate the statistical short-term fairness of SFOS when β is fixed to 50. We set the measurement window length to 100 timeslots, i.e., every 100 timeslots, we count the service discrepancy (in timeslot) received by these two users. In Fig. 3, we first plot the fairness function given by (17) in dotted line, which represents the fairness requirement of the system. Here, we make the weighted service discrepancy, $d_{i,j}$, as the x-axis. Then from (17), the fairness function $g(d_{i,j})$ now has the following expression:

$$g(d_{i,j}) = 1-(1-\exp(-\frac{d_{i,j}}{\beta}\cdot\frac{P_0}{P_i}))\cdot(1-\exp(-\frac{d_{i,j}}{\beta}\cdot\frac{P_0}{P_j})). \quad (23)$$

Following that, we plot the probability achieved by simulation at which the weighted service discrepancy is not less than $d_{i,j}$. We can see the probability is well below the bound given by the fairness function, which evaluates that SFOS can provide excellent statistical fairness to all the flows in the system.

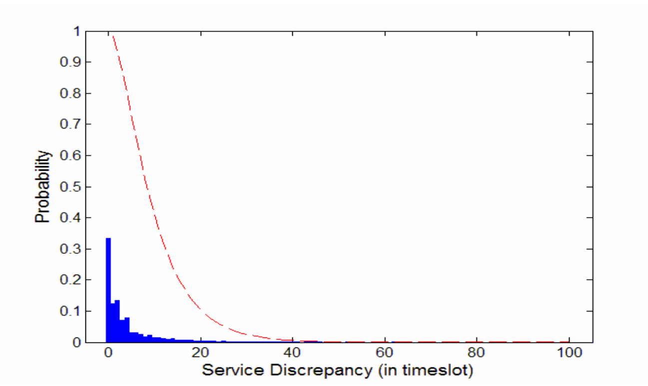

Fig. 3. Statistical short-term fairness guarantee

VI. CONCLUSIONS

This paper introduces SFOS, a wireless scheduling algorithm that provides statistical fairness guarantee over Rayleigh fading channels. Through the design of fairness function and utility function, SFOS allows the system operators have flexibility between perfect temporal fairness and system throughput maximization. Both theoretical analysis and simulation results show that SFOS can obviously improve system throughput while providing excellent fairness.

REFERENCES

[1] F. Berggren and R. Jantti, "Asymptotically fair scheduling in fading channels," IEEE VTC Fall, 2002.

[2] A. Gyasi-Agyei and Seong-Lyun, "Cross-Layer Multiserivce Opportunistic Scheduling for Wireless Networks," IEEE Communications, vol. 44, no.6, pp. 50-57, June 2006.

[3] J.M. Holtzman, "Asymptotic analysis of Proportional Fair algorithm," IEEE PIMRC 2001, San Diego, CA, 2001.

[4] Y. Liu, S. Gruhl, and E. Knightly, "WCFQ: An Opportunistic Wireless Scheduler with Statistical Fairness Bounds," IEEE Trans. Wireless Comm, vol. 2, no.5, pp. 123, Sept. 2003.

[5] P. Goyal, H.M. Vin, and H. Cheng, "Start-time Fair Queuing: A scheduling algorithm for integrated services," ACM SIGCOMM'96, Stanford, CA, 1996.

[6] X. Liu, "Opportunistic Scheduling in Wireless Communication Networks," PhD thesis, Purdue University, 2002.

[7] S.S. Kulkarni and C. Rosenberg, "Opportunistic scheduling policies for wireless systems with short term fairness constraints," GLOBECOM'03, San Francisco, CA, USA, Dec. 2003.

[8] Z. Ji, Y. Yang, J. Zhou, M. Takai, and R. Bagrodia, "Exploiting medium access diversity in rate adaptive wireless. LANs," ACM MobiCom'04, Philadelphia, PA, USA, Sept. 2004.

[9] T.S.E. Ng, I. Stoica, and H. Zhang, "Packet fair queueing algorithms for wireless networks with location-dependent errors," IEEE INFOCOM'98, San Francisco, CA, March 1998.

[10] Y.-C. Wang, Y.-C. Tseng, and W.-T. Chen, "MR-FQ: A Fair Scheduling Algorithm for Wireless Networks with Variable Transmission Rates," Simulation: Transactions of The Society for Modeling and Simulation International, vol. 81, no. 8, pp. 587-608, Aug. 2005.