



Butt, M. M. and Jorswieck, E. A. (2013) Energy Efficient Multiuser Scheduling: Exploiting the Loss Tolerance of the Application. In: 2012 IEEE Global Communications Conference (GLOBECOM), Anaheim, CA, USA, 03-07 Dec 2012, pp. 3531-3536. ISBN 9781467309219.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/155259/>

Deposited on: 15 January 2018

Enlighten – Research publications by members of the University of Glasgow_
<http://eprints.gla.ac.uk>

Energy Efficient Multiuser Scheduling: Exploiting the Loss Tolerance of the Application

M. Majid Butt
Interdisciplinary Center
for Security, Reliability and Trust (SnT)
University of Luxembourg, Luxembourg
Email: majid.butt@uni.lu

Eduard A. Jorswieck
Department of Electrical Engineering and Information Technology
Dresden University of Technology, Germany
Email: eduard.jorswieck@tu-dresden.de

Abstract—We address the problem of energy efficient scheduling for the loss tolerant applications by exploiting the multiuser diversity. The proposed scheduling scheme allows dropping of a certain predefined proportion of data packets on the transmitter side. However, there is a hard constraint on the maximum number of *successively* dropped packets. The scheduler exploits average data loss tolerance to reduce the average system energy expenditure while fulfills the hard constraint on *successively* dropped packets. We analyze the scheme using asymptotically large user limit. The numerical results illustrate the energy efficiency of the scheme as a function of the average packet drop probability and the maximum permitted *successively* dropped packets parameters.

I. INTRODUCTION

Energy efficiency is increasingly getting more attention in the design of modern communication networks. It is important to fulfill the quality of service (QoS) requirements for the users in terms of guaranteed throughput, delay and violation probabilities. However, it is equally important to use every soft requirement on QoS to minimize the cost of transmission.

A lot of literature focuses on energy efficient communication for the delay constrained applications [1], [2]. However, not much work focuses on exploiting the loss tolerance of the application in the scheduling process at the physical layer. It is important to identify the packets which require large energy for transmission and drop them. Though, some researchers have considered similar problems in different settings. Reference [3] discusses a scheduler which differentiates traffic based on the loss and delay tolerance of the user in traditional internet. In [4], the authors address the problem of optimal dropping of packets. They obtain optimal dropping scheme when the size of the packet grows asymptotically large. Reference [5] proposes an algorithm for improving the energy-delay trade-off for the case of dropping a non zero fraction of the packets.

We propose a scheduling scheme which exploits the loss tolerance of the application to improve the system energy efficiency. Loss tolerance is characterized by the average packet loss probability θ_{tar} for a user and her ability to keep the quality of experience acceptable after n or fewer *successively* dropped packets. This constraint on the *successively* dropped packets is referred to as *continuity* constraint in this work. Average packet loss is a degree of freedom that allows the system to drop certain packets to save energy. However, if

more than n packets are dropped *successively*, it becomes difficult to maintain the quality of experience by using complex signal processing and error concealment techniques at the receiver, e.g., Hybrid automatic repeat request (HARQ) scheme with incremental redundancy requires transmission of additional data if a packet is not decoded at the receiver in its first transmission. If the subsequent packet containing the incremental data is dropped to save the energy, it makes the already transmitted packet useless as well and the energy spent on the first transmission gets wasted. Although, higher layer processing can be employed to identify and avoid dropping of such packets, this will not be energy optimal for the resource allocation at physical layer.

Wireless sensor networks (WSN) are an other example of such applications where *successive* dropping of more than a certain number of packets at the sensor nodes may result in inaccurate estimation of the measured data at the fusion node. Similarly, voice and multimedia applications can tolerate loss up to a certain limit but dropping of *successive* packets affects the quality of experience severely. Thus, it is not sufficient to guarantee average packet loss alone. To the best of our knowledge, not much literature deals with the problem of energy efficient scheduling which fulfills both continuity constraint and the average packet loss. In our work, the optimization problem is to choose the set of packets to be dropped such that the constraints on average and *successive* packet dropping are met and the average system energy is minimized.

The rest of this paper is organized as follows. Section II describes the system model. We discuss the proposed scheduling scheme and its large system analysis in Section III and IV, respectively. Special cases for extremely large n and θ_{tar} values are discussed in Section V. We present the numerical results in Section VI and conclude with the main contributions of the work in Section VII.

II. SYSTEM MODEL

We consider a multiple-access system with K users randomly placed within a certain area. Every user is provided an average rate $R = \frac{\Gamma}{K}$ where Γ denotes the spectral efficiency of the system. Γ is normalized by the number of channels M to get spectral efficiency per channel C .

We consider a time-slotted system. Arrivals occur at the start of the time slot; and the scheduling and transmission is completed before the end of the time slot. An uplink scenario is considered.

Each user k experiences a channel gain $h_k(t)$ in slot t . The channel gain $h_k(t)$ is the product of path gain s_k and small-scale fading $f_k(t)$ i.e. $h_k(t) = s_k f_k(t)$. Path loss and small-scale fading are assumed to be independent. The path loss is a function of the distance between the transmitter and the receiver and remains constant within the time scale considered in this work. Small-scale fading depends on the scattering environment. It changes from slot to slot for every user and is independent and identically distributed across both users and slots; but remains constant within each single transmission. This model is often referred to as block fading. For a multi-band system of M channels, small-scale fading over the best channel is represented by, $f_k(t) = \max(f_k^{(1)}(t), f_k^{(2)}(t), \dots, f_k^{(M)}(t))$.

$E_k^R(t)$ and $E_k(t)$ denote the received and the transmitted energy for each user k such that

$$E_k^R(t) = h_k(t)E_k(t). \quad (1)$$

Note that the distribution of $h_k(t)$ differs from user to user. The channel state information is assumed to be known at the transmitter side.

The continuity constraint requires us to allow scheduling of multiple users simultaneously in the same time slot. If only a single user is scheduled per time slot, this constraint cannot be satisfied when multiple users have already dropped n packets.

The scheme follows the results for the asymptotic user case analysis and therefore, there is no limit on the number of users scheduled simultaneously. Those scheduled users are separated by superposition coding. Let \mathcal{K}_m be the set of users to be scheduled in frequency band m . Let $\Phi_k^{(m)}$ be the permutation of the scheduled user indices for frequency band m that sorts the channel gains in increasing order, i.e. $h_{\Phi_1}^{(m)} \leq \dots \leq h_{\Phi_k}^{(m)} \leq \dots \leq h_{\Phi_{|\mathcal{K}_m|}}^{(m)}$. Then, the energy of the user $\Phi_k^{(m)}$ with rate $R_{\Phi_k}^{(m)}$, as scheduled by the scheduler is given by [6], [7]

$$E_{\Phi_k}^{(m)} = \frac{N_0}{h_{\Phi_k}^{(m)}} \left[2^{\sum_{i \leq k} R_{\Phi_i}^{(m)}} - 2^{\sum_{i < k} R_{\Phi_i}^{(m)}} \right]. \quad (2)$$

where N_0 denotes the noise power spectral density. The energy assignment using superposition coding results in the minimum total transmit energy for the scheduled users.

III. MODELING OF THE SCHEDULING SCHEME

We assume that a single packet arrives in each time slot. However, this restriction can be removed by modeling a random arrival process with a constant arrival process with random size [8]. Both of the representations are equivalent as a result of the system level averaging over large number of users in the system. No arrival is modeled by a transmission with zero size and makes no contribution in the system energy. We assume a delay limited system with no delay tolerance where

every packet has to be scheduled or dropped immediately after its arrival. The scheduling decision should take into account the instantaneous channel conditions, average loss tolerance of the application and the history of packet dropping decisions in the previous time slots to avoid dropping a packet if already n packets have been dropped successively. The motivation behind the idea is that we should drop packets in bad channel conditions and exploit good channel conditions for scheduling while respecting the continuity constraint.

We introduce a small-scale fading dependent *dropping threshold* for the scheduling decision of every packet. We use small-scale fading instead of channel gain in the computation of dropping threshold to avoid the phenomenon of *near-far* effect where the users near the base station have better channel gains as compared to the users at the edge of the cell. Scheduling based on small-scale fading gives every user equal opportunities to schedule or drop the packet. If channel conditions are better than the dropping threshold, the packet is scheduled and dropped otherwise. To compute the optimal set of dropping thresholds, we use Markov Decision Process (MDP) theory. We define some terms first.

Definition 1 (State): It is defined as the number of already successively dropped packets for a user k at the arrival time of a packet and denoted by p_k .

As the state depends on the number of successively dropped packets only and the system is completely symmetric with respect to the users with the same state, we drop the subscript k and denote it by p simply.

Definition 2 (Dropping Threshold): It is defined as the minimum fading value required for the transmission of a packet for a user with state p . It is denoted by κ_p . Maximum number of successive packets allowed to be dropped are modeled by n Markov states.

We model and analyze the scheduler in the large user limit. When we have asymptotically large population of the users in the system, the system state model represents the state space of a single user and each user takes her scheduling decisions independent of the other users. Such decoupling principle has been applied in [9], [10] to solve different communication problems.

The state transition mechanism of the scheduler is described as follows. By definition, the number of successively dropped packets determine the state of the process at time t . The state transitions are determined by the state p and the small-scale fading. If a packet is scheduled for transmission, the scheduler returns back to state zero. If it is dropped, the state at time $t + 1$ is $p + 1$. Thus, the state transition probability α_{pq} is written as

$$\alpha_{pq} = \Pr\{S_{t+1} = q | S_t = p\} \quad (3)$$

$$= \begin{cases} 1 - P_f(\kappa_p) & \forall p, q = 0 \\ P_f(\kappa_p) & \forall p, q = p + 1 \\ 0 & \text{else} \end{cases} \quad (4)$$

where $P_f(\cdot)$ denotes the cumulative distribution function (cdf) of the small-scale fading. We are not allowed to drop more

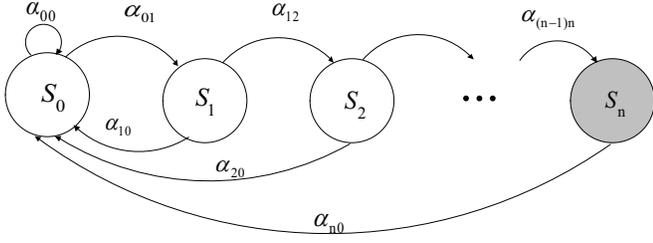


Fig. 1. State diagram for the scheduler with the continuity constraint modeled by a finite state MDP.

than n packets successively. Therefore, κ_n is set to zero to fulfill the condition and terminate the process. The resulting state transition diagram has been shown in Fig. 1.

IV. LARGE SYSTEM ANALYSIS AND OPTIMIZATION

We analyze the proposed scheduling scheme in the large user limit. We treat each transmitted packet as a separate user for the analysis purpose and call it a virtual user (VU). Then, the average energy consumption of the system per transmitted information bit at the large system limit $K \rightarrow \infty$ is given by [7]

$$\left(\frac{E_b}{N_0}\right)_{\text{sys}} = \log(2) \int_0^{\infty} \frac{2^{C P_{h,\text{VU}}(x)}}{x} dP_{h,\text{VU}}(x) \quad (5)$$

where $P_{h,\text{VU}}(\cdot)$ denotes the cdf of the channel gains of the scheduled packets which is composed of the small-scale fading and the path loss.

The transition probability vector $\vec{\alpha} = [\alpha_{00}, \alpha_{10} \dots \alpha_{n0}]$ captures the complete information for the state transition mechanism. The optimization problem is to compute the energy optimal transition probability vector $\vec{\alpha}^*$ which meets the average drop and continuity constraints. The problem is formulated as:

$$\min_{\vec{\alpha} \in \Omega} \left(\frac{E_b}{N_0}\right)_{\text{sys}} \quad (6)$$

$$\text{subject to : } \begin{cases} 0 \leq \alpha_{p0} \leq 1 & p \leq n \\ \theta_r \leq \theta_{\text{tar}} & 0 \leq \alpha_{p0} \leq 1 \end{cases} \quad (7)$$

where Ω denotes the set of permissible vectors for $\vec{\alpha}$ and θ_r is the average dropping probability for a fixed $\vec{\alpha}$ and given by

$$\theta_r = \sum_{p=0}^{n-1} \alpha_{p(p+1)} \pi_p = \sum_{p=0}^{n-1} (1 - \alpha_{p0}) \pi_p. \quad (8)$$

π_p denotes the steady state probability of the state p and follows the property that

$$\sum_{p=0}^n \pi_p = 1. \quad (9)$$

All the forward transitions belong to the events of dropping the packet and the summation over the corresponding transition probabilities $\alpha_{p(p+1)}$ gives the average dropping probability. The corresponding channel-dependent optimal dropping thresholds can be computed from the optimized $\vec{\alpha}^*$ using (4).

We evaluate probability distribution function (pdf) of the channel gain $p_{h,\text{VU}}(x)$ of the scheduled users (packets) using MDP. The scheduling decisions are affected by the small-scale fading distribution only. Thus, the resulting pdf of the small-scale fading of the scheduled VU is given by

$$p_{f,\text{VU}}(y) = \begin{cases} \sum_{p=0}^n c_p \pi_p P_f(y) & y > \kappa_p \\ 0 & \text{else} \end{cases} \quad (10)$$

where $p_f(y)$ and c_p denote the small-scale fading distribution and a constant to normalize the pdf. Equation (10) specifies that a packet is scheduled only if $y > \kappa_p$.

Using (10), the cdf of the scheduled users is given by

$$P_{f,\text{VU}}(y) = \sum_{p=0}^n c_i \pi_p (P_f(y) - P_f(\kappa_p)). \quad (11)$$

The cdf of channel gain $P_{h,\text{VU}}(y)$ is computed¹ using (11) and path loss distribution.

The dropping thresholds are optimized such that the expected system energy cost in each state is minimized. A decision to schedule a packet means that cost of transmission in the current state p is lower than the expected cost of transmission in one of the future $n - p$ states. The constraint on average dropping probability is a soft one and fulfilled on long term basis. The aim of the scheduler is to drop the packets when channel conditions are not good and exploit multiuser diversity to transmit the packets energy efficiently. But the construction of the dropping thresholds should take into account the fact that at least a single packet must be transmitted to avoid successive dropping of more than n packets. If the channel conditions in state n are bad, the event of the forced transmission will be severely suboptimal and its cost will be relatively high. The scheduler aims to compute the balance point where it is beneficial to take the risk of such potentially costly decisions by efficiently designing dropping thresholds depending on the design parameters n and θ_{tar} .

A. Optimization by Simulated Annealing

Optimization of thresholds belongs to a class of stochastic optimization problem where computation of exact solution is quite complex and time consuming, e.g., Traveling salesman problem, Knapsack problem. In literature, many acceptable solutions to such problems have been proposed using algorithms such as genetic algorithm, random search, etc. We employ simulated annealing (SA) algorithm for optimization of thresholds². The advantage of this algorithm is that it accepts a solution with a small probability even if it is worse than the already computed best solution. This step is called muting and helps to avoid local minima. Muting depends on a so called temperature term. At the start of the process, the temperature is very high and muting occurs quite frequently.

¹The closed form expression is evaluated following the techniques presented in [11]. The derivation is omitted here due space limitations.

²The choice is solely based on the wide acceptance of SA to solve such problems. The other techniques like genetic algorithm can be applied as well.

The temperature decreases as the process progresses and so as the muting. This is called cooling. There are many cooling schedules used in literature, e.g., Boltzmann annealing (BA) and Fast annealing (FA) temperature cooling schedules etc. We employ FA in this work [12]. In FA, it is sufficient to decrease the temperature linearly in each step b such that,

$$T_b = \frac{T_0}{c_{sa} * b + 1} \quad (12)$$

where T_0 is a suitable starting temperature and c_{sa} is a constant adjusted according to the requirements of the problem. We skip the details of the algorithm and the interested reader is referred to [8], [13] for the details.

System energy in (5) is the objective function for SA. In our problem, a fixed vector $\vec{\alpha}$ defines a specific configuration. In every iteration, one of the transition probability from the vector $\vec{\alpha}$ is varied and the average packet dropping constraint in (7) is checked. If the constraint is met, the configuration is evaluated using (5). At the end of sufficiently large iterative process, we get a solution $\vec{\alpha}^*$ which is believed to be near optimal in most of the cases.

V. SPECIAL CASES

We investigate the impact of extreme values of the average dropping probability and continuity constraint parameters on the system energy.

A. Extreme Values of Continuity Constraint Parameter

We consider two trivial cases for the extreme values of the continuity constraint parameter n .

- 1) $n = 0$: For this case, the system becomes a lossless system regardless of θ_{tar} and the MDP reduces to a single state with the dropping threshold equals to zero.
- 2) $n = \infty$: It implies that there is no limit on the number of successively dropped packets and average dropping probability is the only constraint to be fulfilled. This condition simplifies the model and the MDP reduces to a single state process where θ_{tar} corresponds to the probability of dropping a packet such that

$$\tilde{\alpha}_{00} = P_f(\kappa_0) = \theta_{tar}. \quad (13)$$

$\tilde{\alpha}_{00}$ denotes the probability of dropping a packet in state zero and κ_0 is the corresponding dropping threshold. If small-scale fading is less than κ_0 , the packet is dropped and scheduled otherwise. The scheduler returns back to state zero regardless of packet scheduling or dropping.

B. Extreme values of Average Dropping Parameter

We specifically treat the cases when target average dropping probability θ_{tar} is relatively large or small as compared to the continuity constraint parameter n . The characterization of these cases is one of the main contributions of this paper.

We first consider the case of large θ_{tar} as compared to the continuity constraint parameter n . To gain motivation for this case, suppose we have no average packet dropping constraint but only continuity constraint. In this case, SA provides us the

*unconstrained*³ (without the average packet drop constraint in (7)) optimized transition probability vector $\vec{\alpha}^*$. We evaluate (8) from $\vec{\alpha}^*$ to get the resulting average dropping probability θ_r^* . The θ_r^* value for this special case is termed as *limiting* average dropping probability and denoted by θ_{lim} . For a fixed n , we will not be able to achieve further energy efficiency by dropping more packets for any $\theta_{tar} > \theta_{lim}$. From these arguments we obtain the following lemma.

Lemma 1: For a fixed continuity constraint parameter n , there exists a finite θ_{lim} such that for all $\theta_{tar} > \theta_{lim}$, the same maximum energy efficiency is achieved as for more restrictive θ_{lim} .

We explain this result with the help of an example. Let us assume a system with $n = 1$ and $M = 1$. The unconstrained optimization gives us an optimized transition probability $\alpha_{01}^* = 0.21$ and the corresponding value for $\theta_{lim} = \theta_r^*$ is obtained from (8). The threshold κ_0 corresponding to α_{01}^* determines the minimum channel conditions where it is better to drop a packet and move to state one. In other words, the cost of transmitting a packet in current channel conditions is larger than the expected cost of transmitting a packet in state one. Remember that a packet *must* be transmitted in state one due to the continuity constraint and there is an associated risk of a forced transmission in a potential bad channel. Thus, the optimized α_{01}^* corresponds to the optimized threshold which allows the user to take this risk.

For this specific example, the intuitive question is why allowing $\theta_{tar} > \theta_{lim}$ does not contribute to the energy efficiency when a user can benefit by dropping more packets? This is explained as follows. The larger θ_{tar} means that the user can drop more packets; but it can drop *only* in state zero as it is prohibitive to do so in state one (because of $n = 1$). When a user drops more packets in state zero and moves into state one frequently, she increases the risk that the channel in state one will be poor. The increased probability of the forced transmissions on (potentially) bad channels in state one enhances the average energy expenditure instead; and energy efficiency is affected adversely. Therefore, when we optimize the thresholds for $\theta_{tar} > \theta_{lim}$ and a fixed n , the optimizer always provides us the set of thresholds obtained for $\theta_{tar} = \theta_{lim}$ and rejects all the $\vec{\alpha}$ after evaluation which result in $\theta_r > \theta_{lim}$.

The energy consumption is a function of the continuity parameter n and the average dropping probability constraint θ_{tar} , i.e., $E_b/N_0(\theta_{tar}, n)$. The larger n and the lower θ_{tar} results in smaller E_b/N_0 . However, at small θ_{tar} the average dropping probability becomes more critical constraint as compared to the continuity constraint. The relationship between target dropping probability θ_{tar} and continuity parameter n is characterized in the following lemma.

Lemma 2: The reduction of the energy consumption E_b/N_0 by increasing the continuity constraint parameter n

³The term unconstrained is only a relative one, referring to the absence of any constraint on average packet drop when n is modeled by the number of states in a finite state MDP. In true mathematical sense, the problem is still constrained due to presence of the continuity constraint.

TABLE I
LIMITING DROPPING PROBABILITY AND SYSTEM ENERGY

n	κ_0	κ_1	κ_2	κ_3	κ_4	$\theta_r^* = \theta_{\text{lim}}$	E_b/N_0
1	0.24	0	-	-	-	0.18	-1.42dB
2	0.54	0.23	0	-	-	0.34	-3.05dB
3	0.76	0.59	0.21	0	-	0.45	-4.07dB
4	0.97	0.79	0.49	0.30	0	0.53	-4.80dB

decreases with decreasing target dropping probability θ_{tar} and becomes negligible at small θ_{tar} .

This result is motivated by the fact that for small dropping probability θ_{tar} , it is hardly feasible to drop more than one packet successively for a user even if the continuity constraint parameter n allows her to do that. For small average dropping probabilities, dropping of more successive packets requires that no (or a few) packet can be dropped any more in the time scale where averaging is performed.

For example, let us assume an averaging window size of 100 and $\theta_{\text{tar}} = 0.02$. This implies that a user can drop packets (on the average) in 2 time slots in a window of 100 time slots. For $n = 1$, the user can drop packets in any two time slots with the worst channels but they *cannot* be adjacent. For $n = 2$, she is allowed to drop the packets even if the time slots with the worst channels happen to be the adjacent ones. But the probability of this event is extremely small and has negligible effects as compared to the system with $n = 1$. A further increase in the value of parameter n would not be beneficial for the system. Both of the lemmas will be further justified by the numerical examples in Section VI.

VI. NUMERICAL RESULTS

We consider a multi-access channel with M bands and assume statistically independent fading on these channels. Every user senses M channels and selects her best channel as a candidate for transmission. The users are placed uniformly at random in a cell except for a forbidden region around the access point of radius $\delta = 0.01$. The path loss exponent α equals 2 and the path loss distribution follows the model described in [7]. All the users experience small-scale fading with exponential distribution with mean one on each of the M channels.

Table I shows the limiting dropping probability and associated system energy for the fixed n when we perform optimization without the average packet drop constraint in (7). The system parameters are $M = 1$ and $C = 0.5$ bits/s/Hz. We will show in the following numerical example that energy efficiency cannot be improved by allowing $\theta_{\text{tar}} > \theta_{\text{lim}}$ for a fixed n .

Fig. 2 shows the optimized system energy for different target average dropping probabilities and continuity constraint parameters. We plot the results for the special case $n = \infty$ for reference. For small θ_{tar} , energy expenditures are almost identical for different values of n and the effect of the continuity constraint is negligible. This confirms Lemma 2 that the continuity constraint parameter n does not have a big impact on the energy efficiency in this region. As θ_{tar} increases, an increase in n does contribute to increase the energy efficiency

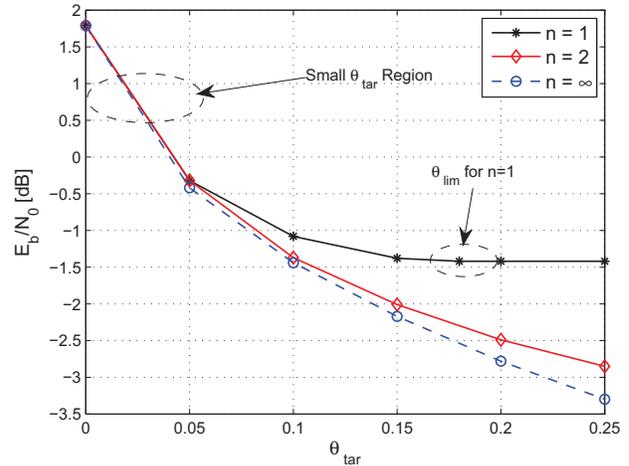


Fig. 2. The system energy as a function of average dropping probability and continuity constraint parameters. The system parameters are $M = 1$ and $C = 0.5$ bits/s/Hz.

of the system. An other important observation is the optimized energy at $\theta_{\text{tar}} > \theta_{\text{lim}}$. From Table I, θ_{lim} equals 0.18 for $n = 1$ and the corresponding energy expenditure is -1.42 dB. We observe in Fig. 2 that the optimization solution provided by SA and the optimized energy remains the same for $\theta_{\text{tar}} = 0.2$ and $\theta_{\text{tar}} = 0.25$. This justifies Lemma 1 that allowing $\theta_{\text{tar}} > \theta_{\text{lim}}$ does not benefit the energy efficiency of the system for a fixed n . Therefore, QoS should always be defined by taking into account both of the parameters jointly.

As we explained in Section IV-A, SA is believed to provide a solution that is acceptable in most of the situations. The SA algorithm uses FA temperature schedule where we simulate 100 temperature values. At each temperature, we evaluate 10 random configurations of every transition probability in $\vec{\alpha}$. For $\theta_{\text{tar}} < \theta_{\text{lim}}$, the system is likely to be more energy efficient if θ_r^* for the solution approaches θ_{tar} closely to benefit from dropping more packets. It does not imply that every configuration which results in larger θ_r is more energy efficient. Some of them are not energy optimal and rejected by the optimizer. Therefore, Δ is only a relative quality measure for the computed solution and defined in terms of θ_r^* for the solution as

$$\Delta = 1 - \frac{\theta_r^*}{\theta_{\text{tar}}} . \quad (14)$$

Δ is expressed in percentage. We observe in Fig. 3 that Δ is quite small for all the cases and the results are acceptable. However, there are more random configurations (transition probabilities) involved in the combinatorial optimization at large n which makes optimization procedure increasingly complex and SA algorithm may require more time to achieve accuracy. For example, Δ is about 13 percent for the case when $n = 4$ and $\theta_{\text{tar}} = 0.1$. When we increase the temperature iterations from 100 to 200, Δ reduces to 4 percent and the optimal system energy decreases correspondingly.

The proposed scheme is based and analyzed in large user limit and holds for large number of users always. However, it

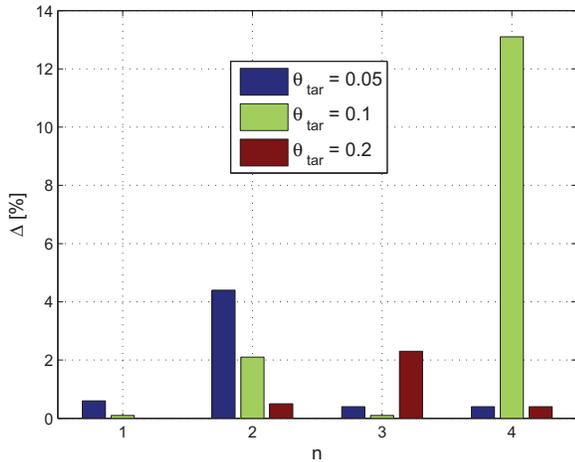


Fig. 3. The accuracy measure Δ as a function of average dropping probability and continuity constraint parameters. The system parameters are $M = 1$ and $C = 0.5$ bits/s/Hz.

is important to show the convergence of the scheme for the finite number of users to claim it to be practical. If the number of users in the system is not large enough, the scheduled sum rate varies greatly from slot to slot and the system energy does not converge to the average system energy. Moreover, the decoupling principle for the users does not hold for the small number of users as explained in Section III. Fig. 4 illustrates the convergence behaviour of the scheme using Monte Carlo simulations. For a fixed number of users and a fixed path loss, we compute variance of system energy for 200 independent small-scale fading realizations. Energy requirement for every scheduled user is computed by (2). Fig. 4 shows that the variance is quite small for a few hundred users and decreases further as the number of users increases. The convergence is faster for the small spectral efficiencies. For the large spectral efficiency values, it requires more scheduled users in each time slot to converge to the average system energy for the asymptotic case. This result provides numerical evidence that the scheme remains practicable for a reasonably small number of users.

VII. CONCLUSION

We proposed a scheduling scheme which exploits the loss tolerance of the application to minimize the system transmit energy over the fading channels. The average and successive packet loss tolerance are the core parameters which jointly define the loss tolerance behaviour of an application. We conclude that average loss tolerance helps to use the radio resources opportunistically and drop the packets which require large expected energy. The successive packet drop parameter poses an additional constraint on resource allocation and needs to be incorporated in the scheduling process. We propose and analyze the scheduler which fulfills both of the constraints and schedule the packets energy efficiently. We show that exploiting loss characteristics of the application for radio resource allocation at physical layer helps to minimize the

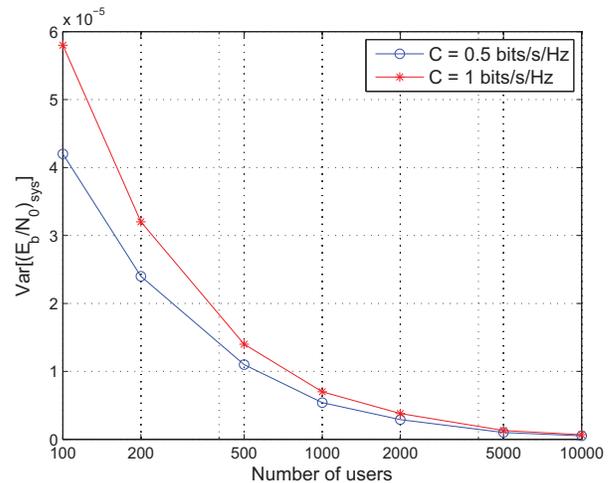


Fig. 4. The finite user convergence of the scheme as a function of number of users for a system with $M = 10$, $n = 1$ and $\theta_{tar} = 0.05$.

system energy expenditures.

ACKNOWLEDGMENT

This work was carried out during an ERCIM "Alain Bensoussan" Fellowship tenure. This Programme is supported by the Marie Curie Co-funding of Regional, National and International Programmes (COFUND) of the European Commission.

REFERENCES

- [1] R. A. Berry and R. G. Gallager, "Communication over fading channels with delay constraints," *IEEE Trans. Inform. Theory*, vol. 48, no. 5, pp. 1135–1149, May 2002.
- [2] T. P. Coleman and M. Medard, "A distributed scheme for achieving energy-delay tradeoffs with multiple service classes over a dynamically varying channel," *IEEE journal on selected areas in communications*, vol. 22, no. 5, pp. 929–941, June 2004.
- [3] A. Striegel and G. Manimaran, "Packet scheduling with delay and loss differentiation," *Computer Communications*, vol. 25, pp. 21–31, 2002.
- [4] W. Chen, U. Mitra, and M. Neely, "Packet dropping algorithms for energy savings," in *IEEE International Symposium on Information Theory (ISIT)*, Nice, France, Jul. 2006, pp. 227–231.
- [5] M. J. Neely, "Intelligent packet dropping for optimal energy-delay tradeoffs in wireless downlinks," *IEEE Trans. on Automatic Control*, vol. 54, no. 3, pp. 565–579, March 2009.
- [6] D. Tse and S. Hanly, "Multi-access fading channels-part I: Polymatroid structure, optimal resource allocation and throughput capacities," *IEEE Trans. Inform. Theory*, vol. 44, no. 7, pp. 2796–2815, November 1998.
- [7] G. Caire, R. Müller, and R. Knopp, "Hard fairness versus proportional fairness in wireless communications: The single-cell case," *IEEE Trans. Inform. Theory*, vol. 53, no. 4, pp. 1366–1385, April 2007.
- [8] M. M. Butt, K. Kansanen, and R. R. Müller, "Hard deadline constrained multiuser scheduling for random arrivals," in *WCNC*, Cancun, Mexico, March 2011.
- [9] M. Benaim and J.-Y. Le Boudec, "A class of mean field interaction models for computer and communication systems," *Performance Evaluation*, vol. 65, no. 11–12, pp. 823–838, 2008.
- [10] P. Viswanath, D. N. Tse, and V. Anantharam, "Asymptotically optimal water-filling in vector multiple-access channels," *IEEE Trans. Inform. Theory*, vol. 47, no. 1, pp. 241–267, January 2001.
- [11] M. M. Butt, "On the near-far gain in opportunistic and cooperative multiuser communications," Norwegian University of Science and Technology, Trondheim, Norway, 2011, Ph.D. dissertation.
- [12] H. Szu and R. Hartley, "Fast simulated annealing," *Physics Letters A*, vol. 122, no. 3, 1987.
- [13] S. Kirkpatrick, C. Gelatt, and M. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, no. 4598, pp. 671–680, May 1983.