

TEMPORAL FAIRNESS ENHANCED SCHEDULING FOR COOPERATIVE RELAYING NETWORKS IN LOW MOBILITY FADING ENVIRONMENTS

Ingmar Hammerström, Jian Zhao, and Armin Wittneben

Swiss Federal Institute of Technology (ETH) Zurich
Communication Technology Laboratory, Sternwartstr. 7, CH-8092 Zurich, Switzerland
Email: hammerstroem@nari.ee.ethz.ch

ABSTRACT

We consider a wireless network, which consists of several active source/destination pairs and a number of idle nodes. Multiple nodes volunteer as relays to facilitate the communication. All nodes have low mobility and the fading channel is essentially constant over the latency time scale of interest (block fading). Conventional adaptive scheduling techniques are either unfair or inefficient in this context. In [1] we proposed a scheduling scheme which does not compromise the quality of service of individual source/destination links (fairness) if the number of relay nodes is sufficiently high. In this paper we propose a simple extension to our scheduling scheme which guarantees short-term temporal fairness among the individual links in the network, while achieving considerable multi-user diversity gains.

1. INTRODUCTION

The use of diversity in the spatial and temporal dimension to mitigate the effects of fading and therefore to increase the reliability of radio links in wireless networks is a well known technique for systems with co-located antennas (space-time coding). Recently a new concept to realize spatial diversity has been introduced in [2] and [3] called *cooperative diversity* or *user cooperation diversity*. The main idea is to use multiple nodes as a virtual macro antenna array, realizing spatial diversity in a distributed fashion. In such a network several nodes serve typically as relays for an active source/destination pair. Relays can be classified as either *decode-and-forward* (DF) or *amplify-and-forward* (AF) relays. AF relays, which are considered in this work, only retransmit an amplified version of their received signals. This leads to low-complexity relay transceivers and lower power consumption since there is no need of signal processing for decoding procedures. Moreover, AF relays are transparent to adaptive modulation techniques which may be employed by the source.

Channel adaptive scheduling is known as an efficient means to achieve multi-user diversity gains in multiple-access networks [4]. It is often referred to as opportunistic scheduling. A good overview about opportunistic scheduling can be found, e.g., in [5–8] and their included references. The authors of [5] present a framework for opportunistic scheduling with different fairness constraints. In [6] and [7] throughput optimal scheduling rules are proposed by taking each users data queue status into account. An opportunistic scheduling extension for IEEE 802.11 systems is presented in [8].

Channel adaptive scheduling is efficient and fair in particular in the high mobility regime, because the channel state varies sufficiently within the latency time scale of interest. Consequently,

each node has a fair chance to have a good link in this time interval. However, fairness is a critical issue particularly in low mobility fading environments. In this case a channel adaptive scheduler, which maximizes the utilization of the shared resources, essentially would only serve the source/destination pair with the best link. At the same time the link throughput of the remaining links would be sacrificed. For quality of service (QoS) reasons, a realistic channel adaptive scheduling scheme in this case has to operate away from the aggregate throughput optimum.

In [9] an *opportunistic beamforming* scheme is proposed to circumvent the drawbacks of low mobility in channel adaptive scheduling in a centralized network. It essentially introduces temporal variations into a quasi-static fading environment by applying time-variant weights at the co-located transmit antenna array. In [10] we propose a simple diversity scheme for distributed cooperative links with AF relays, which translates the block fading time-invariant channel into a time-variant channel by introducing time-variant phase offsets at the relays. The main motivation for the work presented in [1] is the observation, that this time variance could be exploited by considering several source/destination pairs jointly for scheduling purposes. We refer to this approach as *joint cooperative diversity and scheduling* (JCDS). The scheme considerably improves the utilization of the physical resources. It does not compromise the quality of service of individual source/destination links (fairness) if the number of relay nodes is sufficiently high. In this work we extend our JCDS scheme to the case of short-term temporal fairness guarantees [11] which makes fair opportunistic scheduling with a small number of required relays possible.

The remainder of the paper is organized as follows: in section 2 we describe the system model and review our cooperative relaying scheme [10]. In section 3 we summarize our approach of channel adaptive scheduling within a cooperative network. Furthermore, we introduce the extension of our scheduling scheme to the case of temporal share fairness. Performance results are presented in section 4. Conclusions are given in section 5.

Notation: We use bold uppercase letters to denote matrices and bold lowercase letters to denote vectors. Further $(\cdot)^T$, $(\cdot)^*$, $(\cdot)^\dagger$ stand for transpose, conjugate, and conjugate transpose of a matrix, respectively. $\text{diag}[a, \dots, z]$ denotes a diagonal matrix with the elements a, \dots, z on its main diagonal.

2. SYSTEM MODEL AND COOPERATIVE DIVERSITY

In the following we describe the system model for cooperative relaying. The described network consists of N_a source/destination

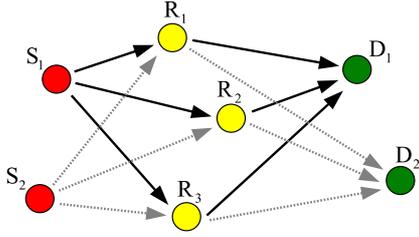


Fig. 1. Two-hop cooperative relaying network consisting of two active source/destination links and three relays. All nodes are single antenna nodes.

pairs and N_r AF relays (e.g., depicted in Fig. 1 for $N_a = 2$ and $N_r = 3$). Communication takes place over two hops (or rather time-slots). In the first slot the source transmits a data packet to the relays and the destination. During the second slot the relays retransmit an amplified version of the received signals to the destination. We assume that all nodes are frame synchronous, i.e., all nodes know the slot timing.

In our scenario mobility of all nodes is low and the channel coefficients are constant over the latency time scale of interest. We assume that the channel is time-invariant over at least one *transmission cycle* (block fading).

To derive an expression for the SNR at the destination the channel coefficients between source and relays (*uplink*) for link n are stacked in the vector $\mathbf{h}_{1,n}$, whereas the channel coefficients between relays and destination (*downlink*) are stacked in $\mathbf{h}_{2,n}$. The direct channel between source and destination is given by $h_{0,n}$. We consider frequency-flat fading, which usually includes path loss, shadowing and small-scale fading. In this paper we assume that all channel coefficients are i.i.d. complex normal random variables $\mathcal{CN}(0, d^{-\alpha})$, where d is the distance between transmitter and receiver and α is the path loss exponent.

Fig. 2 shows the system model of an AF relay. The AWGN contribution for time instance k at relay l is denoted by $m_l^{(k)}$ which is assumed to be $\mathcal{CN}(0, \sigma_R^2)$. Prior to the retransmission the relay amplifies the signal with the factor $g_l^{(k)}$.

To derive the mutual information between the transmitted and the received signals we define the gain matrix $\mathbf{G}_k = \text{diag}[g_1^{(k)}, \dots, g_{N_r}^{(k)}]$. The transmit power of the source is P , and the AWGN contribution at the destination is $\mathcal{CN}(0, \sigma_D^2)$. Neglecting the link index n , the mutual information is then given by

$$I^{(k)} = \frac{1}{2} \log_2 \left[1 + \frac{P|h_0|^2}{\sigma_D^2} + \frac{P|\mathbf{h}_2^T \mathbf{G}_k \mathbf{h}_1|^2}{\sigma_D^2 + \sigma_R^2 \mathbf{h}_2^T \mathbf{G}_k \mathbf{h}_1} \right]. \quad (1)$$

A detailed derivation of the mutual information (1) can be found in [10]. Note that the factor 1/2 accounts for the two channel uses required by the relay traffic pattern.

To achieve a cooperative diversity gain with N_r relays in the block-fading environment it is possible to code the signals at the relays by a multiplication of a data packet with a non-diagonal code matrix [12]. We proposed a simpler approach in [10] which essentially transforms the spatial diversity into temporal diversity by means of a diagonal matrix. This temporal diversity again can be exploited by an outer code introduced at the source. Therefore, our proposed relaying scheme is transparent with respect to adaptive modulation and transmission scheduling.

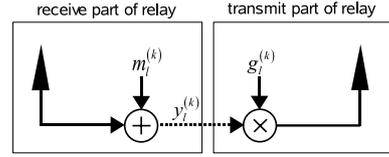


Fig. 2. System model of amplify-and-forward relay

For clarity of exposition we briefly summarize the approach: the linear processing at the relays is time-variant, which results in a time-variant equivalent source/destination channel coefficient (i.e., a time-variant SNR at the destination). In a simple embodiment the relays use a time-invariant gain and a relay-specific time-variant phase offset (phase signature sequence). Thus, the amplification gain of relay l at time instance k is, e.g., given as

$$g_l^{(k)} = \underbrace{\sqrt{\frac{P_R}{P|h_{1,l}|^2 + \sigma_R^2}}}_{\text{time-invariant}} \cdot \underbrace{\exp(j\varphi_l^{(k)})}_{\text{time-variant}}, \quad (2)$$

where $\varphi_l^{(k)}$ denotes the relay-specific and time-dependent phase offset. Each relay transmits with power P_R . We impose a total power constraint P on all relays, i.e., $P_R = P/N_r$. Due to the time-variant phase-shifts at the relays the spatial diversity offered by the N_r relays is translated into the temporal domain. The number of relays limits the maximum achievable diversity order to N_r , because there are only N_r independent downlink channel realizations. To exploit this temporal diversity the input signal sequence is precoded by a linear block code (matrix multiplication) as, e.g., in [13].

Note, that the phases do not necessarily change with every transmitted symbol. We assume that one *transmission cycle* is divided into two time-slots (uplink and downlink) each divided into N_b segments; N_b corresponds to the length of the phase sequences of the relays in the downlink. In this work, the phase sequences are chosen as N_r columns out of an IFFT matrix. But also random phase shifts are appropriate. Note that (2) provide an *instantaneous* relay transmit power P_R and only the phases are time-variant. It would also be possible to additionally model the transmit power as a random process which would ensure only an *average* transmit power of P_R over N_b phase segments.

3. ADAPTIVE SCHEDULING

If we consider a wireless network with multiple source/destination pairs which share the same physical resources the question arises which link should be scheduled. In general there are two extreme optimization criteria. On one hand the scheduling entity wants to maximize the utilization of the physical resources which is usually measured in terms of *system throughput*. On the other hand the fairness among all competing source/destination pairs should be guaranteed which can be measured in terms of *link throughput*. The fairness issue in scheduling again can be divided into two cases: throughput fairness and temporal fairness. In both cases the scheduler gives a certain share of the entire resource to the users. In the first case the shared resource is the whole system throughput. In the latter case, which is considered in this work, the shared resource is time. One scheduling scheme which assures the highest temporal fairness is the static scheduling scheme round robin (RR). The disadvantage of RR scheduling is that it does not incor-

1. Start at beginning of fairness window; reset all counters which represent the number of segments each source/destination pair has been scheduled.
2. Select the pair which supports the highest possible rate and increase counter for this pair.
3. If a scheduled pair reached its maximum temporal share, take it out of the set of pairs which have to be scheduled.
4. Go to next time segment.
5. Repeat 2. - 4. until end of fairness window.

Table 1. Temporal Fairness Enhanced JCDS

porate the actual channel conditions into its decisions. Therefore, it is not possible to achieve multi-user diversity gains.

In the following we shortly review the concept of Joint Cooperative Diversity and Scheduling (JCDS) which was proposed in [1]. It is based on the artificially introduced time-variance of the channel as described in section 2. It makes efficient scheduling in cooperative networks in the low mobility environment possible. In [1] it has been shown that JCDS achieves the same Quality-of-Service (QoS) requirement (defined as the 1%-outage link throughput) as a system with static scheduling. Unfortunately, the number of relays which are required to achieve this is relatively high and increases with the number of source/destination pairs N_a . Therefore, the aim of this work is to extend the JCDS scheme to the case of temporal fairness constraints to reduce the number of required relays. For this, we propose two different approaches which need different amounts of channel knowledge at the scheduling entity.

With respect to our scheduling schemes, we assume that one node of the network serves as central scheduling entity. For illustration we assume that the feedback channel to this node is perfect without any errors.

3.1. Joint Cooperative Diversity and Scheduling

The key idea of the work presented in [1] is to further exploit the artificially introduced time-variance by joint consideration of several source/destination pairs to achieve multi-user diversity. If the fading conditions of all source/destination pairs are symmetric, the proposed scheduler in [1] selects the links on the basis of a greedy (maximum-rate) scheduling metric, i.e., the best link in each phase segment. If the fading conditions are asymmetric the JCDS uses the metric of the proportional fair scheduler [9].

Note that, although fairness of JCDS explicitly benefits from the artificial introduced time-variance it does not assure that each source/destination link is scheduled within one transmission cycle. JCDS only assures temporal fairness over an infinitely large horizon.

3.2. Temporal Fairness Enhanced JCDS

A straightforward extension to short-term temporal fairness for JCDS can be derived as follows. To assure temporal fairness the whole time sequence is divided into fairness windows of length N_f . We count N_f in terms of transmission cycles. $N_f = 1$, e.g., would be the most stringent QoS requirement, because then each pair has to be scheduled within one transmission cycle.

Within this fairness window the scheduler chooses sequentially over the phase segments that source/destination pair which

1. Start at beginning of fairness window; reset all counters which represent the number of segments the source/destination pair has been scheduled.
2. Sort the supported rates for each segment over the source/destination pairs. Calculate the differences of the best and second-best supported rate for each segment.
3. Select the pair and segment which has the highest difference in supported rate; mark segment as scheduled; increase counter for this pair.
4. If a scheduled pair reached its maximum temporal share, take it out of the competition and go back to 2. If not, go back to 3. Repeat until all segments are scheduled.

Table 2. Pre-Selective Temporal Fairness Enhanced JCDS

supports the highest possible throughput. If one pair has reached its maximum temporal share in the actual fairness window it drops out of competition and remains inactive until the next window. We refer to this scheduling scheme as Temporal Fairness Enhanced JCDS, or TF-JCDS. The TF-JCDS scheme is summarized in Table 1.

Obviously, if the first source/destination drops out of competition one order of multi-user diversity is lost, because the scheduler has only $N_a - 1$ pairs left to schedule. From the last remaining source/destination pair it is not possible to achieve any multi-user diversity gain at all. Bearing this matter in mind, it can be seen that the system throughput of TF-JCDS has to be smaller compared to JCDS considered in section 3.1.

3.3. Pre-Selective Temporal Fairness Enhanced JCDS

Another way to guarantee short-term temporal fairness which promises a higher system throughput than the TF-JCDS benefits from the time-invariant physical channel and the deterministic phase sequences at the relays. A scheduling entity which knows the channel coefficients of the network and the phase sequences can determine the supported rates of each source/destination pair in advance for the whole all segments within a transmission cycle. Therefore, it can maximize the system throughput subject to the amount of temporal share of each pair. This optimization problem is hard to solve and needs a lot of computations, in particular, when thinking of a brute force search algorithm. Therefore, we introduce a scheduling algorithm which is optimal for $N_a = 2$ and suboptimal for $N_a > 2$, but achieves at least a higher system throughput as TF-JCDS. We refer to this kind of scheduling scheme as Pre-Selective Temporal Fairness Enhanced JCDS, or PTF-JCDS. It is based on the concept of minimizing the loss in system throughput. This loss arises if a source/destination pair has already reached its maximum amount or temporal share by the sequential scheduling as, e.g. in TF-JCDS, but has higher supported rates to come within the fairness window. Therefore, PTF-JCDS schedules for every segment this pair which shows the highest throughput compared to second best pair. Starting from the segment with the highest difference, this minimizes the loss in system-throughput for $N_a = 2$. The algorithm is summarized in Table 2.

Obviously the required amount of channel feedback to the scheduling entity for PTF-JCDS is much higher than for TF-JCDS. For TF-JCDS only SNR feedback is required which can be ad-

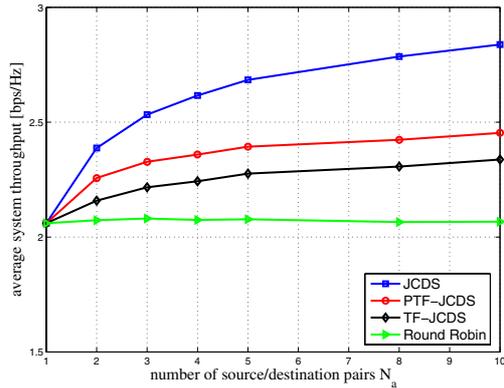


Fig. 3. Average system throughput vs. number of source/destination pairs; $N_r = 10$, $\rho = 10\text{dB}$, $N_f = 1$ transmission cycle.

ditionally much more quantized than feedback of channel coefficients.

Note that for achieving full cooperative diversity on the link level it is necessary that each source/destination pair is scheduled at least N_r times within a transmission cycle. Furthermore, the effective fading channels seen by the link for each scheduled segment have to be uncorrelated, which is only possible if the phase sequences are orthogonal. Therefore, one necessary but not sufficient condition for full diversity regarding the number of segments is: $N_b \geq N_a \cdot N_r$.

4. PERFORMANCE RESULTS

In this section we will present the performance results of our proposed scheduling schemes by means of computer simulations.

Simulation Setup: We consider a network where N_r relays are placed at random uniformly distributed on a disk with diameter d_0 . The N_a sources and destinations are placed on the border of this disk, each on opposite sides. The distance between all nodes on the border of the disk are chosen to be maximum. Therefore, the angle between two subsequent sources (or destinations) is chosen as π/N_a . We assume that all channel coefficients are i.i.d. complex normal random variables $\mathcal{CN}(0, d^{-\alpha})$, where d is the distance between transmitter and receiver and $\alpha = 3$ the path loss exponent. Furthermore, we assume that the noise variance of the relays is equal to the noise variance of each destination, i.e. $\sigma_R^2 = \sigma_D^2 = \sigma^2$. The parameter $\rho = P/(\sigma^2 \cdot d_0^\alpha)$ denotes the average SNR of the direct link between source and destination.

Note, that because the source/destination pairs are placed on the border of the disk nearly symmetrical fading conditions can be assumed for each pair, i.e., the SNR distributions at the destinations are the same for all pairs. Nevertheless, we use the proportional fair scheduling metric [9] for the JCDS scheme.

Results: In Fig. 3 the average system throughput of the proposed scheduling schemes vs. the number of source/destination pairs N_a is shown for $N_r = 10$ relays and a SNR of $\rho = 10\text{dB}$. The length of the fairness window is set to $N_f = 1$ transmission cycle, while the number of segments per transmission cycle is set to $N_b = 400$.

The 1%-outage link throughput for this set of parameters is depicted in Fig. 4. The 1%-outage link throughput is the throughput each individual link would have in 99% of all cases within the fairness constraint window of length N_f . It is directly related to QoS parameters like data rate and delay.

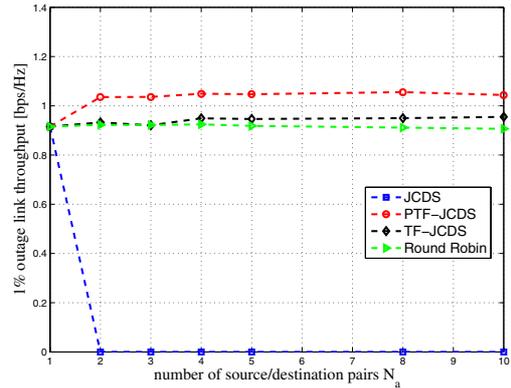


Fig. 4. 1%-outage link throughput vs. number of source/destination pairs; $N_r = 10$, $\rho = 10\text{dB}$, $N_f = 1$ transmission cycle.

The multi-user diversity gain with increasing number of source/destination pairs is clearly visible in Fig. 3 for the opportunistic scheduling schemes. The system throughput of the RR scheme is independent of the number of source/destination pairs. The JCDS without temporal fairness constraints shows the highest average system throughput. But regarding the link throughput in Fig. 4 it shows the worst performance. It is not able to assure a nonzero 1%-outage link throughput within $N_f = 1$ transmission cycle. We only introduced phase shifts at the relays and set an instantaneous power constraint on the relay transmit power P_R . Nevertheless the introduced time-variance is not sufficient to enable the JCDS to have a nonzero 1%-outage link throughput. To increase the time-variance, introducing time-variant relay transmit power could be a choice (compare to last paragraph in section 2).

The TF-JCDS achieves approximately the same outage link throughput as the static scheduling scheme RR, because the scheduler operates sequentially and does not predict when a certain source/destination pair has its maximum. Therefore, it is likely that the remaining pair (all others have already reached their maximum allowed temporal share) has to be scheduled without regarding its actual supported rate. Note that this performance is a big improvement compared to JCDS.

The PTF-JCDS has the global knowledge about all channel coefficients and is therefore able to determine when a certain source/destination pair is at maximum or minimum in advance and can therefore achieve higher outage link throughput compared to RR and TF-JCDS, while achieving also the highest multi-user diversity gain with subject to temporal fairness constraints.

In Fig. 5 we relax the stringent condition on the temporal fairness window size $N_f = 1$ transmission cycle. We compare the average system throughput over an increased length of N_f (in terms of transmission cycles). It is shown for $N_a = 10$ (solid lines) and $N_a = 2$ (dashed lines) source/destination pairs. It can be seen that both schemes, the TF-JCDS and the PTF-JCDS approach to the performance of the JCDS without temporal fairness constraints quite fast.

Note, that the performance of the PTF-JCDS is based on perfect prediction of the variations of the physical channel which changes for every transmission cycle. In a real system this performance would depend on the temporal correlation of the channel coefficients. The TF-JCDS on the other hand does not need this prediction and is therefore a good scheduling scheme with lower feedback requirements compared to PTF-JCDS. It has the same

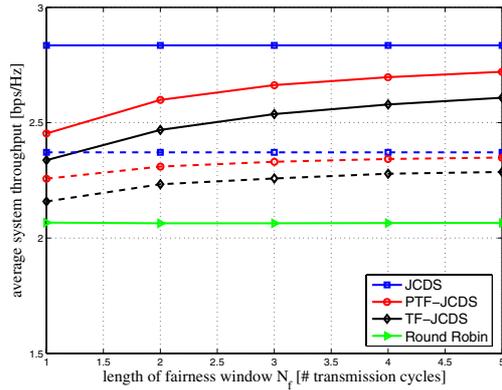


Fig. 5. Average system throughput vs. length of fairness window N_f for $N_a = 10$ (solid lines) and $N_a = 2$ (dashed lines); $N_r = 10$, $\rho = 10$ dB.

feedback signaling requirements as JCDS without short-term fairness constraints. Therefore, from the feedback point of view there is a trade-off between temporal fairness assurance with high QoS requirements and maximum system throughput scheduling with only long-term temporal fairness.

In Fig. 6 the 1%-outage link throughput vs. length of the fairness window N_f is depicted. It is shown for $N_a = 10$ (solid lines) and $N_a = 2$ (dashed lines) source/destination pairs. It can be seen that the outage link throughput within the fairness constraint window also increases, which is mainly due to the increased temporal diversity by the physical channel. The TF-JCDS still performs the same as the static RR in terms of QoS and is independent of the number of active source/destination pairs. The PTF-JCDS profits most from the increased length of the fairness constraint window N_f . Particularly, it can be seen that for $N_a = 2$ source/destination pairs also the JCDS without short-term temporal fairness assurance achieves a nonzero outage link throughput for $N_f > 2$ transmission cycles.

5. CONCLUSIONS AND OUTLOOK

We extended our in [1] presented JCDS to the case of short-term fairness guarantees. We proposed two different approaches: the TF-JCDS which only needs SNR feedback and the PTF-JCDS which needs full channel knowledge. We conclude, that both approaches achieve at least the same QoS as a RR scheduler while achieving considerable multi-user diversity gains. Future work will be concerned with the signaling which is required for scheduling and with the verification of our results in our RACooN Laboratory [14].

6. REFERENCES

- [1] A. Wittneben, I. Hammerstroem, and M. Kuhn, "Joint cooperative diversity and scheduling in low mobility wireless networks," in *Proc. GLOBECOM*, (Dallas, TX), Nov. 29 – Dec. 3, 2004.
- [2] N. J. Laneman, D. N. Tse, and G. W. Wornell, "Cooperative diversity in wireless networks: Efficient protocols and outage behavior," *IEEE Trans. Inform. Theory*, vol. 50, pp. 3062–3080, Dec. 2004.
- [3] A. Sendonaris, E. Erkip, and B. Aazhang, "User cooperation diversity-Part I: System description," *IEEE Trans. Commun.*, vol. 51, pp. 1927–1938, Nov. 2003.

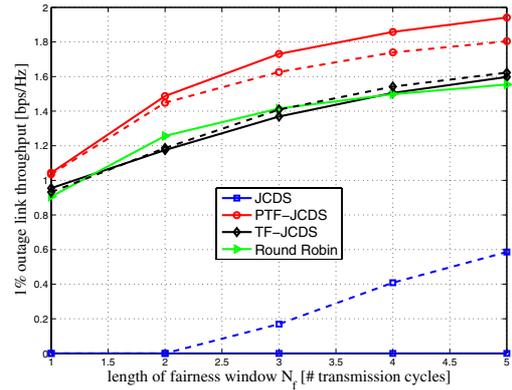


Fig. 6. 1%-outage link throughput vs. length of fairness window N_f for $N_a = 10$ (solid lines) and $N_a = 2$ (dashed lines); $N_r = 10$, $\rho = 10$ dB.

- [4] P. Bhagwat, P. Bhattacharya, A. Krishna, and S. Tripathi, "Enhancing throughput over wireless LANs using channel state dependent packet scheduling," in *Proc. IEEE INFOCOM '96*, (San Francisco, CA), pp. 1133–1140, Mar. 24–28, 1996.
- [5] X. Liu, E. Chong, and N. Shroff, "A framework for opportunistic scheduling in wireless networks," *Computer Networks*, vol. 41, pp. 451–474, Mar. 2003.
- [6] M. Andrew, K. Kumaran, K. Ramanan, A. Stolyar, P. Whiting, and R. Vijayakumar, "Providing quality of service over a shared wireless link," *IEEE Commun. Mag.*, vol. 39, pp. 150–154, Feb. 2001.
- [7] S. Shakkottai and A. Stolyar, "Scheduling for multiple flows sharing a time-varying channel: The exponential rule," tech. rep., Bell Laboratories, Dec. 2000.
- [8] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proc. MOBICOM*, (Atlanta, GA), Sep. 23–28, 2002.
- [9] P. Viswanath, D. N. C. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inform. Theory*, vol. 48, pp. 1277–1294, June 2002.
- [10] I. Hammerstroem, M. Kuhn, and A. Wittneben, "Cooperative diversity by relay phase rotations in block fading environments," in *Proc. SPAWC*, (Lisbon, Portugal), Jul. 11–14, 2004.
- [11] S. S. Kulkarni and C. Rosenberg, "Opportunistic scheduling policies for wireless systems with short term fairness constraints," in *Proc. GLOBECOM*, vol. 1, (San Francisco, CA), pp. 533–537, Dec. 1–5, 2003.
- [12] Y. Jing and B. Hassibi, "Distributed space-time coding in wireless relay networks-Part I: Basic diversity results," *IEEE Trans. Wirel. Comm.*, July 2004. (Submitted for publication).
- [13] Y. Xin, Z. Wang, and G. B. Giannakis, "Space-time diversity systems based on linear constellation precoding," *IEEE Trans. Wirel. Comm.*, vol. 2, pp. 294–309, Mar. 2003.
- [14] "RACooN Laboratory of the Swiss Federal Institute of Technology." www.nari.ee.ethz.ch/wireless/research/projects/racoon/introduction.html.