Considerations on Forwarder Selection for Opportunistic Protocols in Wireless Networks

Anatolij Zubow, Mathias Kurth, and Jens-Peter Redlich

Humboldt University, Germany.
E-mail: {zubow, kurth, jpr}@informatik.hu-berlin.de

***Abstract*—Opportunistic Routing gained lots of attention as a way to improve the performance of wireless multi-hop relay mesh networks. The key characteristic is its ability to take advantage of the numerous, yet unreliable wireless links in the network. The most important part of every opportunistic routing protocol is the forwarder selection algorithm. Most of the currently known protocols assume that signal paths from a sender to all candidates are independent to each other and therefore resulting in independent packet error rates. However, this assumption does not hold for spatially close candidates which are not so uncommon in indoor networks. In this paper, we present empirical measurements from our 802.11 indoor test-bed which reveal that signal paths to spatially close nodes are correlated. We believe that the loss in the radio propagation due to shadow fading is similar to spatially close nodes. For our setup we find out that a spatial correlation exists when the nodes are closer than 2m to each other. We present a candidate set selection algorithm which is able to calculate the packet error rate of a candidate set even when the individual packet error rates are correlated, e.g. due to spatial correlation. Therefore only a simple modification to the existing ordinary link probing is required. Finally, we present modifications we made to our packet-level simulator to respect spatial correlation as well as simulation results.**

***Index Terms*—Wireless Mesh Networks, Opportunistic Routing, Measurements**

I. INTRODUCTION

Wireless multi-hop relay networks are not only a popular research area; today they also play an important role for sensor networks as well as community networks [11]. In contrast to fixed networks, a well-known problem in wireless networks is the existence of channel fading, causing intermittent connectivity or changing link capacity. A very promising wireless communication idea that inherently considers radio aspects is the notion of multi-user diversity [12]. The idea here is to communicate with users at good instances, generally interpreted as exploiting channel variations and then selecting the best user(s). The diversity idea, through user data broadcasting in high node density networks, has since then been enhanced in the Extremely Opportunistic Routing (ExOR) [3], [4]. The main idea of opportunistic routing is that, instead of selecting a single node to be the next-hop forwarder for a packet, multiple nodes can potentially act as the next-hop forwarder. An important part of every opportunistic routing protocol is the forwarder selection algorithm. Most of the currently known protocols assume that signal paths from a sender to all candidates are independent to each other and therefore the packet error rate (PER) of a candidate set can be computed easily from the particular PER's. However, this assumption does not hold for spatially close candidates which are not so uncommon in indoor networks.

In this paper, we present empirical measurements from our 802.11 indoor test-bed which reveal that signal paths to spatially close nodes are correlated resulting in dependent PER's from the sender towards these nodes. We believe that the loss in the radio propagation due to shadow fading [14] is similar to spatially close nodes. Signals coming from the same direction experience a similar shadow effect and the location dependency of the shadow fading values makes sure that another node at the same location also gets the same signal attenuation [8]. To our best knowledge, no work was done in investigating if such a spatial correlation of the shadow fading values also exists in 802.11 indoor networks. We will present measurements which will indicate that such an effect also exists in indoor wireless networks and propose a decorrelation length of $d_{corr} = 2m$ for the Gudmundson autocorrelation function [10].

The main contributions of this paper include: (1) empirical measurements from an 802.11 indoor test-bed that reveal that signal paths to spatially close nodes are correlated if the nodes are less than 2m away from each other, (2) a candidate set selection algorithm which is able to calculate the PER of a candidate set even when the individual PER's are correlated, e.g. due to spatial correlation, (3) by extending the ordinary link probing correlations between different signal paths can be efficiently calculated, (4) we show how the model of the wireless channel within a packet level simulator can be improved by assuming that the shadowing fading values of spatially close nodes are correlated. Here we also present simulation results.

II. RELATED WORK

The extreme opportunistic routing (ExOR) was proposed by Biswas et al. in [3], [4] for wireless multi-hop networks. In order to skip intermediate hops towards the destination, ExOR selects a group of candidate nodes as the potential next-hop forwarder. The candidate set contains the best next hops (smallest ETX (Expected number of Transmissions) value [5]) of every path towards the destination in terms of ascending path length. The node with the smallest ETX value to the destination that successfully receives the current packet will relay the packet further. The selection of a forwarder and the prevention of unnecessary transmissions rely on the slotted

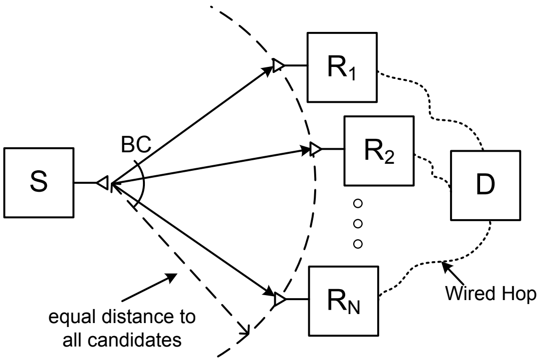

Figure 1. Experimental setup. The source node S simulates an opportunistic transmission by broadcasting data packets towards candidates $R_1 - R_N$.

acknowledgment, a distributed agreement among candidates and sender. Candidates coordinate the transmission of the acknowledgement according to their priority, i.e. their position in the candidate set. Among all candidates that successfully received the frame, the node with the highest priority sends the first acknowledgement and is responsible for relaying the packet. However, there is no guarantee that all candidates agree on the same forwarder. We already discussed several scenarios in which missed acknowledgements result in redundant transmissions and duplicates [19]. Several approaches addressing this problem have been proposed: batch processing [4], robust (i.e. passive, piggybacked) acknowledgement [17] and compressed slotted acknowledgement [19]. Many variations of opportunistic routing protocols have also been proposed [15], [18], [20]. All of the above work is trying to find a good metric for forwarder selection. There is a tradeoff between large candidate sets and overhead due to distributed agreement between them (acknowledgements) as well as unnecessary transmissions (e.g. duplicates). To avoid duplicates most approaches demand that the candidates are well connected to each other. However, this can lead to candidates which are spatially very close to each other. Further, all presented approaches assume that the signal paths from the sender to all potential forwarders are uncorrelated. With measurements in our 802.11 test-bed we will show that this assumption does not necessarily hold for dense networks (here indoor), where it is not unusual that nodes are spatially-close to each other. However, if candidates with correlated PERs' are chosen the opportunistic gain is significantly reduced.

A very important aspect of radio propagation is the loss due to shadow fading [9], [14]. This is caused by the presence of obstacles which lie in the propagation path of the radio waves. While modeling the shadow fading loss, the spatial, angular and temporal correlation of the shadow fading values has to be considered [8]. Gudmundson et al. proposed a one-dimensional model for the spatial correlation properties for suburban and urban environments [10]. The normalized autocorrelation function R(Δx), where Δx is the change in distance, can be described with sufficient accuracy by an exponential function as

$$R(\Delta x) = exp(-\frac{\Delta x}{d_{corr}} ln2) \tag{1}$$

with the decorrelation length d_{corr}, which is dependent on the environment. For the urban vehicular test environment (VTE), [7] proposes d_{corr} = 20m. Further, the author claimed that the correlation works for distances up to 500m. The authors of [8] presented two-dimensional shadow fading profiles for outdoor UMTS/GSM networks which can be used in simulations of realistic mobile radio network scenarios. We believe that the correlated shadow fading of spatially close nodes is the reason why we observe a correlation of link PER's from the sender to its candidates. For our indoor 802.11 test-bed we propose a value of d_{corr} = 2m.

Parameter	**Value**
Physical layer	802.11b
Link layer	broadcast
RF channel	1, (2412MHz)
Bit rate	1Mbit/s
No. receivers	3
UDP payload size	1500 bytes
UDP flow duration	25sec
No. runs	10

Table I
PARAMETERS USED DURING MEASUREMENTS IN 802.11 TEST-BED.

III. MEASUREMENTS

In order to validate our assumption that the signal paths from a sender to multiple spatially close receiver nodes are correlated and therefore cannot be assumed to be independent we conducted experiments in our 802.11 indoor test-bed. The test-bed spans three houses on four different floors of our institute's building. The building mainly consists of concrete and drywalls as well as glass claddings. There are many other IEEE 802.11b/g Wireless LAN access points in the building. Soekris routers equipped with an Atheros AR5213a-based 802.11b/g WiFi card were used. All nodes are connected to our wired LAN to allow monitoring without interfering with the node's wireless operations. For the following experiment a sender was chosen in the third floor, whereas the receivers were placed in the fourth floor. During the measurement the distance between the receiver nodes where varied. However, we tried to keep the distance between the sender and each receiver constantly at around 20m. The experiment setup is depicted in Fig. 1. The sender S is sending broadcast packets at the maximum transmission rate. Each by a receiver (R_1-R_3) successfully received packet is annotated with the receiver's address and forwarded via the wired backbone to the final destination D. The destination keeps statistics about which packet was received by which node. The information about which packet was received by which node is used to calculate the correlation of the received packets between different candidate nodes. In the following we are considering the setup consisting of one sender and three receivers (R_1-R_3). During the experiment the distance between the receivers was varied from 5cm to 500cm. Here the case with the spatially close nodes (5-15cm) is very typical for receivers equipped with multiple radios. For the measurement the following parameters where used (Table I).

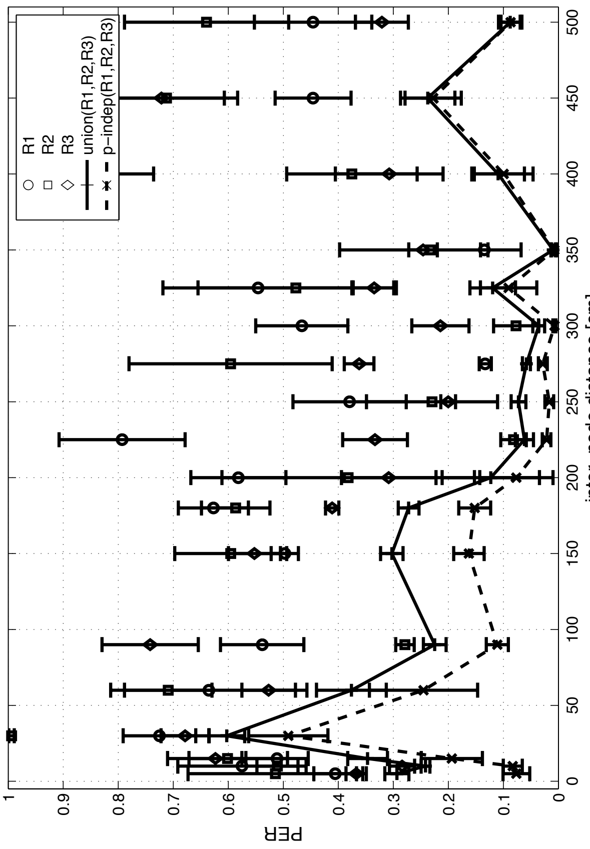

Figure 2. PER in dependence of the distance between the receivers: (a) for each receiver alone (R1-R3), (b) if the received packets of the three nodes are combined (union(R1,R2,R3)), (c) the calculated PER when we assume that the individual PERs' are independent (p-indep(R1,R2,R3)).

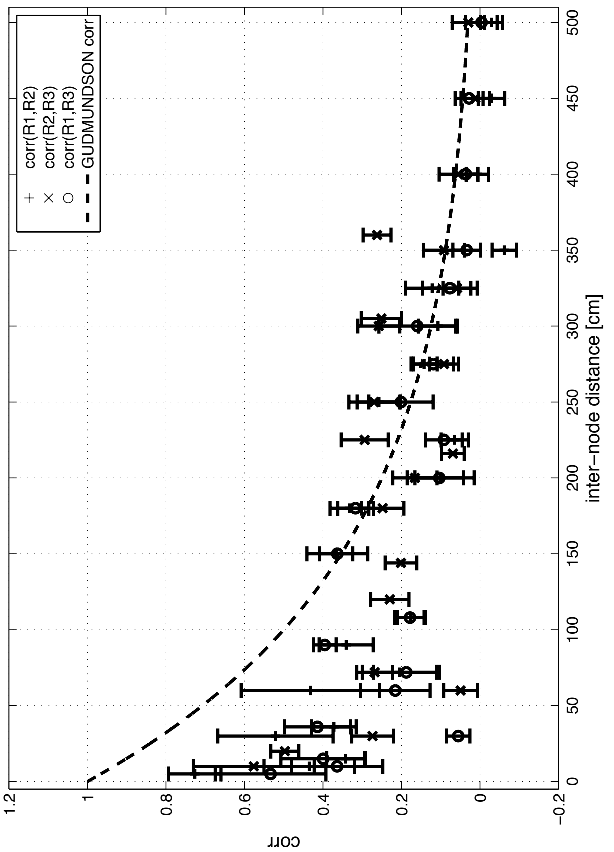

Figure 3. Correlation of received packets between different receivers (corr) in dependence of the distance between the receivers.

A. Results

Due to the high overhead of coordination between the candidates most opportunistic protocols (e.g. ExOR) perform selection-combining on packet-level. That is, an opportunistic transmission is successful if the packet is successfully received by at least one candidate. This can be seen as a form of macro-diversity. In contrast, most multi-antenna systems perform combination of the different signal paths on symbol-level, which is called micro-diversity. From the literature we know that multiple antennas at the receivers can provide sufficient antenna-diversity if the antennas are at least spaced by half the wavelength. Now the question is if this also holds for macro-diversity approaches. In other words: what is the minimum spatial distance between the candidates so that the signal paths on macro-level (i.e. packet) become uncorrelated?

Fig. 2 shows the PER at the receivers (R_1-R_3). First, one can see that most of the links offer a PER different from 0. Furthermore, most of the links are variable in time (high standard deviation). Finally, the distance between the receiver nodes has only a limited influence on the link quality between the sender and each receiver. Now, imagine the candidate set consists of the receivers (R_1,R_2,R_3). An opportunistic transmission fails if none of the candidates successfully receives the packet (union(R_1,R_2,R_3)). As you can see this dramatically reduces the PER especially for nodes which are far away from each other. Consider as an example the case d=500cm. The mean PER's of the individual receivers R_1-R_3 are 0.45, 0.64 and 0.32. The PER of the combined receiver (union(R_1,R_2,R_3)) is 0.08 - combining three moderate links results in a link with a low effective PER. The advantage is smaller for spatially close receivers. Consider as an example the case d=5cm. Again, the mean PER's of the individual receivers R_1-R_3 are 0.4, 0.51 and 0.37, which are comparable to the former case. However, the PER of the combined receiver remains high (0.3). The reason for that is that the signal paths are no longer uncorrelated if the receivers are placed spatially close to each other. This results in correlated PER values.

In Fig. 3 the correlation between the received packets at each receiver pair is depicted. The correlation value is calculated as follows: At first for each node we calculated a bit vector representing the received (1) and not received packets (0). E.g. the bit vector contains a 0 at position x if the packet with number x was not successfully received by that node. Now corr(R1,R2) represents the correlation between the bit vectors of node R1 and R2. One can see that receivers which are close to each other have higher correlation values than receivers which were placed far apart. That means that if the receiver are close to each other (corr = 0.5) and one of them does not successfully receive a packet it is very likely that the other candidate also misses that packet. This reduces

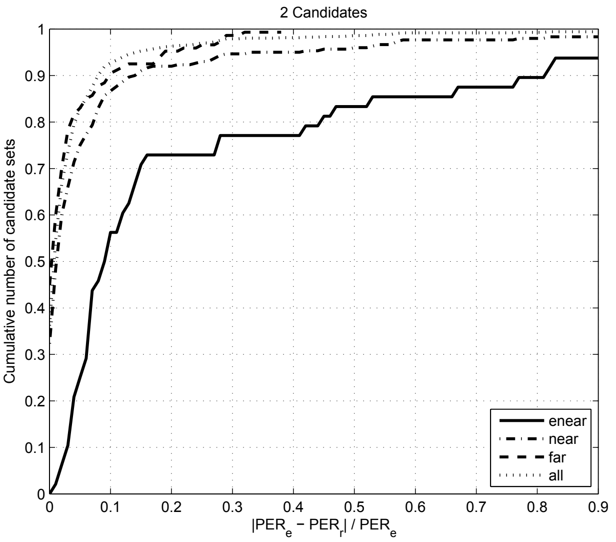

Figure 4. Cumulative number of candidate sets as a function of the error between the real (PER_r) and expected (PER_e) PER value for a candidate set size of 2.

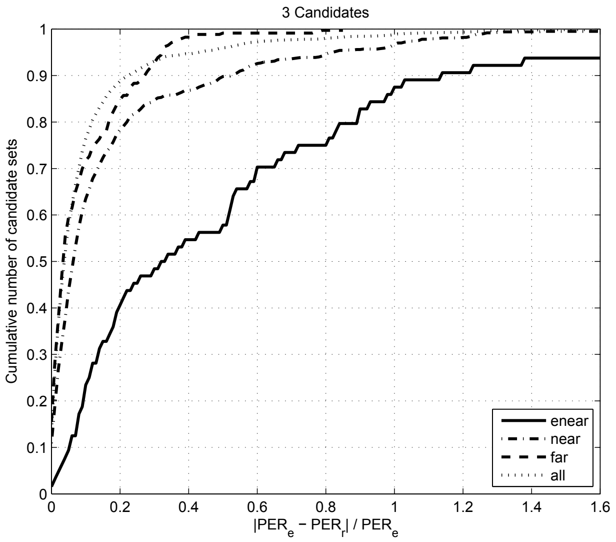

Figure 5. Cumulative number of candidate sets as a function of the error between the real (PER_r) and expected (PER_e) PER value for a candidate set size of 3.

the advantage of having the second candidate. By increasing the distance between the receiver nodes the correlation values decreases. For the case where the distance is around 500cm the receivers are uncorrelated. If we say that a correlation value of greater than 0.3 means that the receivers are correlated then only receivers which are more than 2m away from each other can be handled as independent. For all shorter distances this is not the case. Given our results from Fig. 3 we can make a rough estimation for the d_{corr} in equation (1). For our setup a good value seems to be around 1m.

B. Discussion

Most forward selection algorithms assume that the signal paths to receiver are independent to each other. Therefore, they are choosing candidates in a way that the product of their individual packet error probabilities is minimized. Furthermore, to avoid duplicates they choose candidates which are well connected to each other. The latter one leads to candidates which are spatially close to each other. To show the problem of such an approach, we also computed the PER of the candidate set under the assumption that the signal paths of all candidates are fully uncorrelated. The value is depicted in Fig. 2 as p-indep(R_1,R_2,R_3). The relative difference between this value and the real PER value can be considered as an error. This value is very large especially for spatially close nodes ($< 2m$). This error is depicted in Fig. 4 and 5 for a candidate set of a size of 2 and 3.

At first we consider the case with 2 candidates (Fig. 4). In the case of spatially very close nodes ($\leq 15cm$) there are 22% candidates sets where the relative difference between the expected and the real PER is greater than 40%. If the distance between the nodes is smaller than 100cm for 5% of the candidate sets the error remains. Only if we consider all node placements ($\leq 500cm$) the error is around 8% and occurs only for 5% of the candidate sets. This error grows with the size of the candidate set. With 3 candidates (Fig. 5) and spatially very close nodes nearly every second candidate set offers an error of more than 50%. This error also remains high for setups where the nodes are at most 100cm away from each other. If we consider all nodes the error is around 8% for 10% of the candidate sets.

IV. CANDIDATE SELECTION

Most opportunistic protocols use link probing for neighbor discovery and link quality estimation (loss rate and capacity). Every node periodically broadcasts link probes on a single or multiple bit-rates. On receiving a link probe each node updates its neighbor list accordingly. Furthermore, it maintains a history of link probe receptions to calculate the individual link costs (e.g. ETX or Expected Transmission Time (ETT, [6])). Within a link probe packet the calculated link cost of all neighbors are locally distributed so that every node knows the link costs to neighboring nodes within the hop count distance of 2. Finally, the path metric of a multi-hop route is calculated by summing up the individual link costs along the path.

In the following we will present two proposals for the calculation of the best candidate set. The first approach assumes that there is no correlation in the PER towards the candidates, whereas the second algorithm also considers correlated PER values.

A. Basic Algorithm

Let $N = \{n_1, .., n_M\}$ be the set of nodes in the network. Assume that a node $s \in N$ needs to send a packet towards destination $d \in N$ in a multi-hop fashion. The task of the candidate set selection algorithm is to choose the best candidates for the next hop towards d. The best candidate set is the one with the largest expected forwarding progress towards the final destination in terms of the path metric:

$$\begin{aligned} \mathrm{cssa}(s,d) &= \{c \mid c \in \mathrm{pcs}(s,d) \wedge \forall \bar{c} \in \mathrm{pcs}(s,d) : \\ &\qquad m(s,d,c) \geq m(s,d,\bar{c})\} \end{aligned} \tag{2}$$

where $\mathrm{pcs}(s,d)$ represents the set of potential candidate sets for a transmission from s to d. Here we are only considering candidate sets with a size of up to $\mathit{max\_cs}$ nodes:

$$\mathrm{pcs}(s,d) = \{cs \mid cs \in \mathcal{P}(\mathrm{nb}(s,d)) \wedge |cs| \leq \mathit{max\_cs}\} \tag{3}$$

where $\mathcal{P}(.)$ represents the power set and $\mathrm{nb}(s,d)$ is the set of all neighbors of s that have a path metric towards the destination d that is not greater than that of node s:

$$\mathrm{nb}(s,d) = \{n \mid n \in N \wedge \mathrm{pm}(n,d) \leq \mathrm{pm}(s,d)\} \tag{4}$$

here $\mathrm{pm}(s,d)$ represents the path metric from node s to d. Finally, we define function m which assigns a numerical value to a given candidate set. A higher value represents a more favorable candidate set - higher reception probability and/or greater forwarding progress towards the final destination:

$$\mathrm{m}(s,d,c=(c_1,..,c_S)) = \sum_{i=1}^{S} p_{\mathrm{fwd}}(s,c,i) \times \mathrm{pg}(s,d,c_i)\} \tag{5}$$

where $\mathrm{pg}(s,d,c_i)$ represents the forwarding progress towards the final destination in terms of the path metric. Candidates that are closer to the final destination are more favorable:

$$\mathrm{pg}(s,d,c) = \mathrm{pm}(s,d) - \mathrm{pm}(c,d)\} \tag{6}$$

Finally, $p_{\mathrm{fwd}}(s, c = (c_1,..,c_S), i)$ represents the probability that the candidate c_i from the candidate set c will be the next forwarder of the data packet. In the simplest case we can assume that the distributed coordination (slotted acknowledgment) among the candidates is perfect, i.e. the candidate that successfully received the data packet and is closest (in terms of the path metric) to the final destination will be the next forwarder. We assume that the candidates $c = (c_1,..,c_S)$ are sorted in ascending order according to their path metric to the final destination; i.e. candidate with a lower index are closer to the destination:

$$p_{\mathrm{fwd}}(s,c=(c_1,..,c_S),i) = p(s,c_i) \times \prod_{j=1}^{i-1}(1-p(s,c_j)) \tag{7}$$

where $p(s,c_i)$ represents the probability that the data packet send by s will be correctly received by candidate c_i.

According to the basic algorithm the estimation of $p(s,c_i)$ is simple. Since the algorithm assumes that the PER's towards the candidates are independent, it can simply use the loss rate obtained from ordinary link probing. Consider the following example: node s calculates the metric m for the candidate set $cs = (y,x)$ as follows. The PER obtained from link probing is $p(s,x) = 0.5$ and $p(s,y) = 0.5$. Let y be the highest candidate and the forwarding progress of x and y is $\mathrm{pg}(s,d,x) = 100$ and $\mathrm{pg}(s,d,y) = 200$ respectively. Node y is the next forwarder if the packet is received by y. It is irrelevant if the packet is also received by x since y is the highest candidate. Furthermore, node x is the next forwarder if and only if the packet is received by x and not by y. Based on this information we are able to calculate the metric for the candidate set: $\mathrm{m}(s,d,c=(x,y)) = 0.5 \times 200 + 0.5^2 \times 100 = 125$.

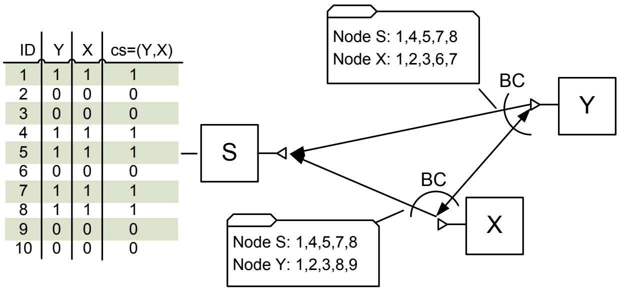

ID	Y	X	cs=(Y,X)
1	1	1	1
2	0	0	0
3	0	0	0
4	1	1	1
5	1	1	1
6	0	0	0
7	1	1	1
8	1	1	1
9	0	0	0
10	0	0	0

Figure 6. With the help of the IDs enlisted in the link probes a node can estimate the PER towards a candidate set.

B. *Advanced Algorithm*

The advanced algorithm extends the basic algorithm by assuming that the PER's towards the candidates are correlated. To be able to calculate the PER of a given candidate set we need to extend the link probing in the following way. The link probes are not only used to calculate the loss rate of individual links, but also that of candidate sets. The empty payload in the link probe packets is used to inform neighboring nodes about which particular link probe the node receives from its neighbors. Each link probe packet is identified by a unique number. With the help of this information a node can calculate the PER for each candidate set. Consider our example illustrated in Fig. 6. Node s is able to calculate not only the particular PER's towards x and y, but also the PER of the candidate set $cs = (y,x)$. Node s calculates the metric m for the candidate set $cs = (x,y)$ as follows. Assume that the packet reception at x fully depends on the reception at y as depicted in Fig. 6, i.e. if a packet is successfully received by y than the packet is also received by x with probability of 1. Note, that the PER's from link probing remain the same. The metric for the candidate set is: $\mathrm{m}(s,d,c=(x,y)) = 0.5 \times 200 + 0 \times 100 = 100$. Here, we can see that the metric is lower (100 vs. 125) than the one calculated by the basic algorithm. In general the basic algorithm tends to overestimate the metric.

V. SIMULATIONS

Based on our observations we made in the real test-bed we improved our packet-level simulator. At first we present our model of the wireless channel. Here we assume that the shadowing fading of spatially close nodes is correlated and only the small-scale fading is independent. Based on this model we conducted several simulations which show that in networks with high spatial correlation (dcorr) the opportunistic advantage decreases significantly. However, the performance can be increased when using our proposed algorithm from section IV-B.

A. *Wireless Channel*

The most important characterization of the mobile wireless channel is the variation of its signal strength over time and frequency [16]. We can identify three pair wise independent multiplicative effects, which have an impact on the radio propagation: large-scale path loss, shadowing and multi-path fading [1]. We used the following radio chain within our simulator. At first the signal is reduced by large-scale fading. Thereafter the signal is further reduced by medium-scale fading via the shadowing model [14]. Channel impairments due to multi-path are calculated with the help of the small-scale fading. Finally, the signal is exposed to noise. This is modeled with an AWGN channel. Finally we are assuming a block fading model, i.e. the fading is constant for the duration of a packet transmission. We assume that the shadowing fading of spatially close nodes is correlated and only the small-scale fading is independent. The shadowing fading is usually described by means of a lognormal distribution. The pdf of the SNR γ is given by:

$$\begin{aligned} f_\gamma(\gamma;\sigma,\bar{\gamma}) &= \frac{1}{\sqrt[2]{2\pi}\cdot ln(\sigma)\cdot\gamma} \\ &\cdot \quad \exp\Big(\frac{-(\ln(\gamma)-ln(\bar{\gamma}))^2}{2\cdot\ln^2(\sigma)}\Big) \qquad (8) \end{aligned}$$

where $\bar{\gamma}$ is the mean SNR according to the large-scale path loss and σ is a variable parameter.

Imagine a network with n nodes. In most simulators (e.g. NS2) the shadowing is realized by a standard Gaussian random vector w which is a collection of n i.i.d. standard Gaussian random variables w_1, ..., w_n. Since we assume that the shadow fading of spatially close nodes is no longer independent we used the algorithm of Park et al. to generate correlated random variables [13]. The algorithm is able to generate correlated Poisson random variables. However, a Poisson distribution with parameter λ is approximately normal for large λ. The approximating normal distribution has parameters $\gamma = \sigma^2 = \lambda$. Finally, we calculated the correlation matrix with the help of equation (1) and the geographical positions of the nodes. The decorrelation length d_{corr}, which dependents on the environment was varied throughout the simulations. Fig. 7 shows an example of the shadow fading experienced at different receivers. A small value for d_{corr} results in small correlation between the shadowing values (upper plot). This changes if we use a high value for d_{corr} (lower plot). Here the shadowing values are similar.

B. Performance Evaluation

JiST/SWANS is used in our simulation study [2]. We extended the simulator to support the 802.11a specification. The simulation was run on the topology depicted in Fig. 1. The parameters we used are summarized in Table II. The simulation results are depicted in Fig. 8. We can make two observations. First, if we assume a spatial correlation of the shadowing fading with a value of $dcorr = 10$ then the opportunistic advantage decreases significantly. In case of $dcorr = 10m$ and a candidate set of a size of 2, 3 and 4 the throughput decreases by 7-14%, 10-22% and 15-26% compared to the result we get when using a value of $dcorr = 1m$, respectively. For higher $dcorr$ values the decrease is becoming even more dramatic. However, we can improve the throughput by using our algorithm from section IV-B (Fig.9). The throughput increases by 9-21%, 10-33% and 13-38% for a candidate set of a size of 2, 3 and 4 respectively.

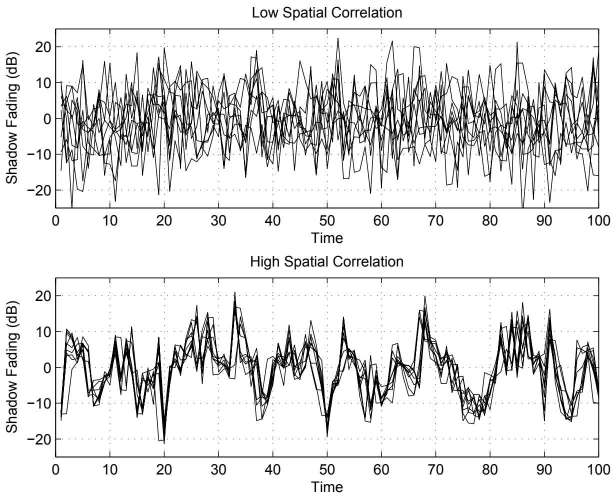

Figure 7. Shadowing fading observed by spatially-close receivers. Each line represents a different receiver. With low spatial correlation each receiver may experience a totally different value (upper figure). However, with higher spatial correlation coefficient this becomes unlikely (lower figure).

Simulation Parameter	Value
Propagation model	Shadowing,Rayleigh
Path loss exponent β	3.5
Shadowing standard deviation σ	$8dB$
Decorrelation length d_{corr}	$1, 10m$
Physical layer	802.11g (6Mbit/s)
No. nodes	20
Placement model	Random (radial)
No. candidates	$1-4$
UDP flow duration	$10sec$
Seeds	$\#100$

Table II
SIMULATION PARAMETERS.

VI. CONCLUSION

With the help of measurements in an 802.11 indoor-test-bed we investigated the effect of the spatial distance between

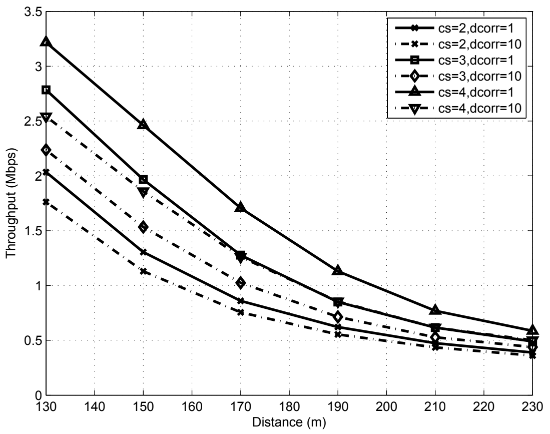

Figure 8. By increasing the decorrelation length d_{corr} the throughput of an opportunistic transmission decreases.

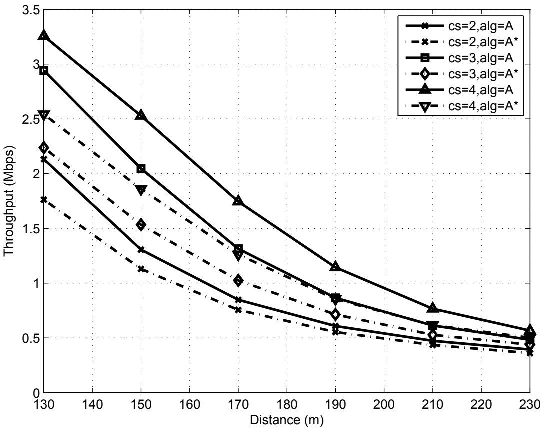

Figure 9. By considering that the PER's towards the candidates are correlated the advanced candidate selection algorithm (A*) is able to outperform the basic algorithm (A). A decorrelation distance of $dcorr = 10$ was used.

candidates in a forwarder set of an opportunistic transmission on the PER. We found a spatial correlation between nodes which are closer than 2 m from each other in our 802.11 indoor scenarios. We presented a candidate set selection algorithm which is able to calculate the PER of a candidate set even when the individual PER's are correlated. Finally, we presented modifications we made to our packet-level simulator to respect spatial correlation as well as simulation results.

In our ongoing research we are systematically evaluating different scenarios (indoor as well as outdoor) to find confident values which can be used to improve simulation models. Furthermore, we are trying to find out whether in 802.11 networks also an angular and temporal correlation of the shadow fading values exists.

REFERENCES

[1] BALANIS, C. A. *Antenna Theory. Analysis and Design.* Wiley & Sons, 2005.

[2] BARR, R., HAAS, Z. J., AND VAN RENESSE, R. JiST: an efficient approach to simulation using virtual machines. *Software: Practice and Experience 35*, 6 (2005), 539–576.

[3] BISWAS, S., AND MORRIS, R. Opportunistic routing in multi-hop wireless networks. *SIGCOMM Comput. Commun. Rev. 34*, 1 (2004), 69–74.

[4] BISWAS, S., AND MORRIS, R. ExOR: opportunistic multi-hop routing for wireless networks. In *Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)* (2005), vol. 35, pp. 133–144.

[5] DE COUTO, D. S. J., AGUAYO, D., BICKET, J., AND MORRIS, R. A high-throughput path metric for multi-hop wireless routing. In *International Conference on Mobile Computing and Networking (MobiCom)* (2003), pp. 134–146.

[6] DRAVES, R., PADHYE, J., AND ZILL, B. Routing in multi-radio, multi-hop wireless mesh networks. In *International Conference on Mobile Computing and Networking (MobiCom)* (2004), pp. 114–128.

[7] ETSI. Tr 101 112, selection procedures for the choice of radio transmission technologies of the umts (umts 30.03). *Tech. Rep., European Telecommunications Standards Institute* (1998).

[8] FORKEL, I., SCHINNENBURG, M., AND ANG, M. Generation of two-dimensional correlated shadowing for mobile radio network simulation.

[9] GOLDSMITH, A. *Wireless Communications.* Cambridge University Press, 2005.

[10] GUDMUNDSON, M. Correlation model for shadow fading in mobile radio systems. *Eltronics Letters 27* (1991), 2145–2146.

[11] HUMBOLDT UNIVERSITY BERLIN, SYSTEMS ARCHITECTURE GROUP. Berlin Roof Net project. www.berlinroofnet.de, 2007.

[12] KNOPP, R., AND HUMBLET, P. A. Information capacity and power control in single-cell multiuser communications. *IEEE International Conference on Communications (ICC)* (1995), 331–335.

[13] PARK, C.G., P. T., AND SHIN, D. A simple method for generating correlated binary variates. *The American Statistician 50* (1996), 306–310.

[14] RAPPAPORT, T. S. *Wireless Communications: Principles and Practice (2nd Edition).* Prentice Hall PTR, 2001.

[15] ROZNER, E., SESHADRI, J., MEHTA, Y., AND QIU, L. Simple opportunistic routing protocol for wireless mesh networks. *2nd IEEE Workshop on Wireless Mesh Networks* (2006).

[16] TSE, D., AND VISWANATH, P. *Fundamentals of Wireless Communications.* Cambridge, 2005.

[17] ZHONG, Z., AND NELAKUDITI, S. On the efficacy of opportunistic routing. In *IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)* (2007).

[18] ZHONG, Z., WANG, J., AND NELAKUDITI, S. Opportunistic any-path for-warding in multi-hop wireless mesh networks. *USC CSE Technical Report* (2006).

[19] ZUBOW, A., KURTH, M., AND REDLICH, J.-P. Multi-channel opportunistic routing in multi-hop wireless networks. Tech. rep., Humboldt University Berlin, 2006.

[20] ZUBOW, A., KURTH, M., AND REDLICH, J.-P. Multi-channel opportunistic routing. In *European Wireless Conference (EW)* (2007).