Interferer Nulling Based on Neighborhood Communication Patterns

Robert Vilzmann[1], Jörg Widmer[2], Imad Aad[2], and Christian Hartmann[1]

[1] Technische Universität München, Institute of Communication Networks,
Munich, Germany `lastname@tum.de`
[2] DoCoMo Euro-Labs, Future Networking Lab, Munich, Germany
`lastname@docomolab-euro.com`

1 Introduction

Smart antennas have a number of useful properties that help to improve the performance of wireless devices. Beamforming allows to emit power in a given sector around a device, by phase shifting the signals sent by different antenna elements so that they constructively add in a specific direction. Similarly, antenna patterns can be modified to have a very small antenna gain in the direction of interferers or source of noise (nulling). Prior work has shown that smart antennas positively affect network properties such as connectivity, path length, interference, etc. in comparison to omni-directional antennas [1, 2]. The increase in hardware complexity for smart antennas and the increase in protocol complexity to exploit their properties vary significantly depending on the specific solution, but smart antennas have even been used successfully for very resource constrained devices such as sensor nodes [3].

Existing approaches can be divided into two categories. In work that falls into the first category, the beamforming direction is either fixed [3] or random [2]. Slightly more sophisticated algorithms may select the beamforming direction based on aggregate information collected from the neighborhood, e.g., setting the beamforming direction based on the number of neighbors found in specific directions [4]. This results in low MAC protocol complexity and a relatively stable topology, since the beamforming direction changes slowly (if at all). Higher performance gains are possible with schemes that fall into the second category: communication-based beamforming. Here, the beamforming direction may be changed on a per packet basis. This approach, however, substantially increases MAC layer and routing complexity, having to deal with directional deafness and very quick changes in the communication topology that may be hard to control. It also requires that nodes frequently infer the direction in which the communication partner lies through Direction of Arrival (DoA) estimation [5], or have very accurate information about the positions of their neighbors [6, 7].

In this paper we explore an alternative solution intending to reduce interference and increase the spatial reuse of radio resources. It can be used without prior beamforming, but also on top of fixed or random direction beamforming. It's complexity is still limited compared to communication-based beamforming on a per packet basis. The proposed algorithms adapt the beamforming pattern to the communication pattern in a given neighborhood. This communication pattern is inferred from the forwarding table of a node, as well as the overheard traffic. We investigate a simple algorithm where nodes start out with an omni-directional antenna pattern. A node distinguishes between neighbors with which it communicates and "undesired" neighbors by looking at the local next hop information. When the node overhears a certain number of packets from a neighbor with which it does not communicate, the node places a "null" in the antenna pattern in the direction of this neighbor. To allow establishment of new routes between nodes where a null exists, it is possible to either periodically send and receive control packets in an omni-directional manner, or let the nulls time out after a certain period.

This paper is organized as follows. Sec. 2 explains the proposed nulling approach in more detail. Modeling assumptions and the considered beamforming method are summarized in Sec. 3. Simulation results are shown and discussed in Sec. 4 and Sec. 5. We conclude with Sec. 6.

2 Algorithms

The goal of the algorithms is to adapt the shape of a node's antenna pattern in order to increase the SINR of communicating nodes. To this end, antenna nulls can be placed toward other nodes to attenuate interference. Furthermore, placing the antenna main lobe toward communicating nodes can enhance the signal reception. Nulling can be combined with both directional and omni-directional antenna patterns. This work will, however, only consider antenna nulls on top of omni-directional patterns.

For a given node, a neighbor is called a *desired* node if the node either transmits to or receives from that neighbor in connection with any of the data flows in the network. All nodes that are not desired nodes are called *undesired*.

2.1 The Basic Approach

Each node in the network annotates its neighborhood table with information from the PHY layer about signal strength. Neighbors are divided into *desired* (communicating nodes) and *undesired* (potential interferers), based on the forwarding table and previous communication events. In order to reduce the interference and thus increase the SINR, a node places a null towards the undesired neighbor with the highest signal strength. With a reasonable number of antenna elements, the remaining degrees of freedom of the antenna array can then be used to maintain the original (omni-directional) pattern as much

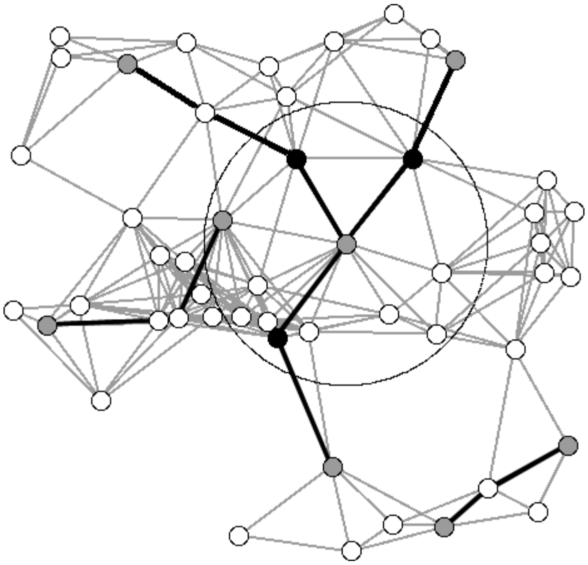

(a) Before nulling. Dark lines indicate data flow routes.

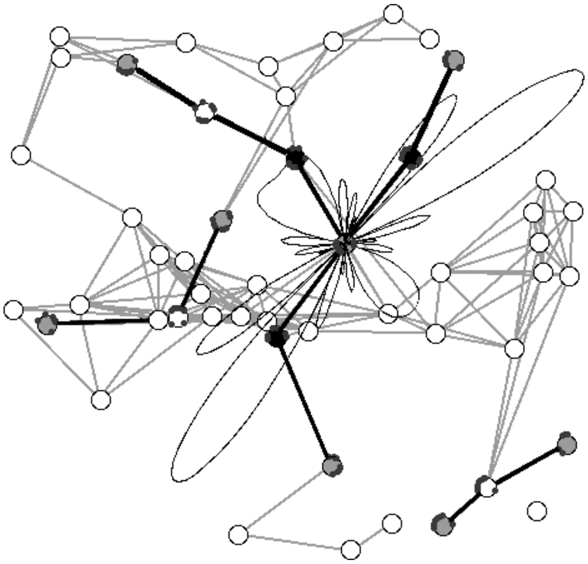

(b) After nulling. Small gray dots indicated antenna null directions.

Fig. 1. Sample scenario with four data flows (gray nodes are end points). For the node under consideration (center of the indicated antenna pattern), the four data flows result in three desired nodes (black). All remaining nodes (gray and white) are undesired nodes, in case they are close enough to be detected.

as possible. No further topological dynamics are thus introduced, and desired communication links can be preserved.

2.2 Advanced Nulling Approaches

Placing a null in one interferer's direction is usually sub-optimal and a node may place further nulls (on several interferers) to increase its signal quality. However, in practice, placing a null in a given direction also changes the shape of the whole antenna pattern. Therefore, placing a null towards an interferer does not necessarily increase the SINR, since it may decrease the antenna gain toward desired nodes.

Finding the optimal solution requires a node to check all combinations of placing nulls toward interferers to attenuate their signals without decreasing the gain of the antenna main lobes, to keep a satisfactory SINR level to communicating nodes. Due to the complexity of such an approach, we adopt a simple greedy algorithm here. It may also be possible to use different beamforming techniques than the one described in Sec. 3.3. However, we use the latter in this work since it has desirable properties from a networking perspective, such as explicit null directions and transparent behavior of the antenna response, relatively low complexity, and applicability to both omni-directional patterns and gain-maximizing beamforming prior to placing antenna gain nulls.

In our greedy approach as used in the simulations section, a node first starts establishing a list of undesired and desired neighbors, and nulls the one with the highest interference power. The node then checks whether the link budget of desired neighbors is still satisfactory (i.e., no connections to communicating nodes are lost so as to maintain the existing routes), due to the possible changes in the antenna beam pattern. If the signal quality degrades below a given threshold the node removes the null. Irrespective of whether or not the null is placed, the node subsequently places a null toward the (second) next highest interferer, and the same procedure is repeated. The algorithm ends after nulling all undesired neighbors in the list, or upon reaching the nulling limitations at the node, which depend on the number of antenna elements.

An example nulling scenario is shown in Fig. 1. After nulling as many undesired nodes as possible, the links toward the desired nodes are still present. While all active nodes carry out the nulling procedure, only the pattern of the sample node is illustrated for the sake of clarity. Small gray dots indicate null directions.

2.3 Updating Antenna Nulls

Several factors mandate recomputing null directions:

- Changes in network topology: when old communicating nodes (resp. interferers) move out of and new ones move into a node's neighborhood.

- Changes in communication pattern: when communication sessions end and communicating nodes become potential interferers (and vice versa).
- Changes in channel conditions: when fading and shadowing may change the signal strength from communicating nodes and interferers.

The frequency of recomputing null-directions depends on the application scenario and must take the above factors into consideration.

3 Antenna and Network Modeling

In this work we apply a nulling technique to a realistic antenna model, and estimate to what extent interference can be eliminated by applying nulling on the physical layer. The corresponding modeling assumptions are summarized in this section, including the channel and link model.

3.1 Channel and Link Model

In order to model the wireless channel, we apply the modified free-space model with an attenuation exponent α. By applying this model, we implicitly assume line-of-sight propagation. That is, no explicit multipath propagation is considered. With this model we can describe the received power p_r at a node located at distance s from a sending node transmitting with power p_t as:

$$p_r = p_t \cdot g_t \cdot g_r \cdot \left(\frac{\lambda}{4\pi s}\right)^2 \left(\frac{s_0}{s}\right)^{\alpha-2} , \tag{1}$$

where λ is the carrier wavelength, s_0 is a reference distance, g_t is the antenna gain of the transmitter in direction toward the receiver, and g_r is the antenna gain of the receiver in direction toward the transmitter. The directive antenna gains at the nodes, which are effective for transmission as well as for reception, are given by the beam pattern of the adaptive antennas. How these beam patterns are generated is explained in Sec. 3.3.

Now that we have the channel model in place we can describe the link model, i.e. we can determine when two nodes can establish a communication link between each other. For this purpose we define a threshold power value p_{r0}, referred to as the receiver sensitivity. If the received power p_r according to (1) is not lower than the sensitivity p_{r0}, we assume that a communication link can be established. According to the channel model and considering that the antenna gain patterns are effective both for transmission and reception, links are always bi-directional. Note that this does not guarantee that data packets are always received successfully on such a link. Rather, packets can still be lost when packet collisions (severe interference from parallel transmissions in the neighborhood) occur.

In addition to the receiver sensitivity p_{r0}, we apply another power threshold, the detection threshold p_{rd}: If the received power p_r according to (1) is lower than the receiver sensitivity p_{r0} but not lower than the detection threshold p_{rd}, the nodes cannot establish a communication link, since the received power would be too weak for decoding. However, the signal can still be detected and the received power is assumed to be sufficiently high in order to perform *Direction of Arrival* (DoA) estimation. Hence, in such a scenario, nodes can detect each other as interferers and are able to estimate the direction from which the interference is coming, which can then be used to advise the beamformer to produce an antenna pattern that eliminates this interference.

The actual detection threshold used in our simulations ($p_{rd} = -91\,\mathrm{dBm}$) together with the applied noise figure (temperature 295 K, bandwidth 22 MHz, noise factor 1.0) yields an SNR threshold of 9.5 dB, which should be sufficient for DoA estimation (cf. e.g. [8]), assuming that interference from other nodes does not harm the DoA estimation.

3.2 Antenna Model

We assume that nodes are equipped with m-element *Uniform Linear Array* (ULA) antennas. That is, m antenna elements are placed in one line with a fixed distance $\Delta = \lambda/2$ between neighboring elements, where λ represents the carrier wavelength. Each element is modeled as an isotropic radiator. The overall transmit power of the array is fixed to p_t, however the power can be distributed unevenly among the antenna elements, such that $\sum_1^m p_{t,i} = p_t$, where $p_{t,i}$ represents the power radiated by the i-th element. The power-to-element assignment is determined by the antenna steering vector $\mathbf{a} = (\mathbf{a_1}, \ldots, \mathbf{a_m})^{\mathbf{T}}$. The steering vector comprises m complex factors, determining the amplitude and phase shift for each antenna element. The actual beamforming method, i.e. the choice of the steering vector depending on the requested number and directions of nulls in the pattern is explained in Sec. 3.3. Throughout this paper we assume ULA10 antennas, i.e. $m = 10$ antenna elements.

3.3 Nulling Method

Our goal is to investigate the potential of improving the performance of a multi-hop network by nulling out interferers. However, nulling out interferers requires changes in the beam pattern of nodes, which in turn can change the connectivity of the network significantly. Obviously, we would like to be able to control these changes in topology in the following way:

- The link to the undesired interferer should be removed, while
- the links to the other neighbors in the connectivity graph should remain. This is to avoid rapid changes in the topology and thus in the routing tables.

In order to approach these goals, we apply a beamforming method which places nulls in a desired direction while minimizing the deviation (in terms of the mean square difference) of the new pattern from the pattern before the null was added. In particular, we adopt the beamforming method described in [9]. It is briefly sketched in the following.

Since we are assuming that the antenna elements comprising the array are isotropic radiators, the antenna pattern can be represented by the *array factor*. In the case of an ULA with a spacing of $\lambda/2$ which we are considering here, we can express the array factor $\kappa(\theta)$ as a function of the angle $0 \leq \theta < 2\pi$

$$\kappa(\theta) = \sum_{n=1}^{m} a_n e^{-jn\pi \sin\theta} \ , \tag{2}$$

when the steering vector $\mathbf{a}$ is applied. Assume the current array factor is $\kappa_0(\theta)$, determined by the steering vector $\mathbf{a_0}$. This pattern might either be omnidirectional or it might already have some desired nulls at angles $\theta_1, ..., \theta_{\mu-1}$. Now we want to place an additional null in the pattern at angle θ_μ. Since we want the new pattern to be as similar as possible to $\kappa_0(\theta)$, we have to choose the new steering vector $\mathbf{a_a}$, yielding the new array factor $\kappa_a(\theta)$, which solves the following optimization problem:

$$\min \left\{ \epsilon(\kappa_a) = \int_0^{2\pi} |\kappa_0(\theta) - \kappa_a(\theta)|^2 d\theta \right\} \ , \tag{3}$$

subject to

$$\kappa_a(\theta_i) = 0 \ , \qquad i = 1, \ldots, \mu \ . \tag{4}$$

This problem can be solved in close form as detailed in [9], as long as $\mu < m$. That means, we can place up to $m - 1$ distinct nulls in our pattern. Exemplifying beamforming patterns for the considered antenna configuration are illustrated in Figs. 2, 3, 4. Note that the antenna pattern of linear arrays is always symmetric. Each antenna gain null implies a null in the direction symmetric to the array.

3.4 Network and Traffic Model

For our simulations we generate random scenarios in which 50 nodes are independently placed according to a uniform distribution in an area of 500×500 m^2. For each scenario multi-hop traffic is generated. Active nodes, i.e. nodes generating traffic, randomly choose one of the 49 remaining nodes as sink, and determine the shortest path to this sink in terms of the number of hops. We will look at two different setups: in the first setup, 10 nodes out of the 50 nodes are active. In the second, all 50 nodes are active and generate traffic. For all figures in this paper, statistics were averaged over 1000 random networks, and all contemplable nodes in these networks.

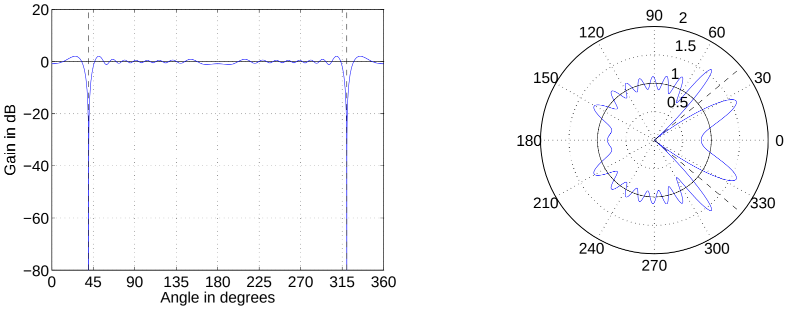

(a) Attenuation/amplification in dB. (b) Linear antenna gain.

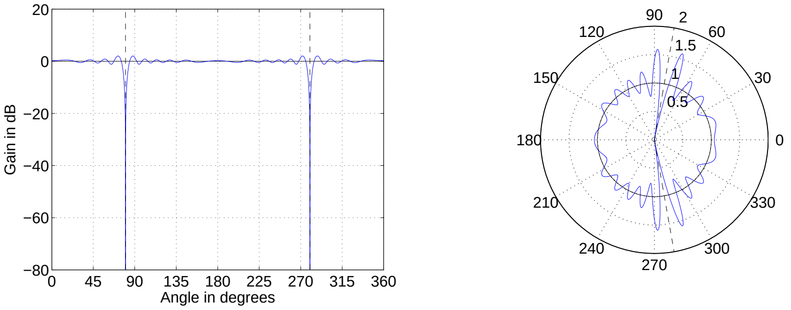

(c) Attenuation/amplification in dB. (d) Linear antenna gain.

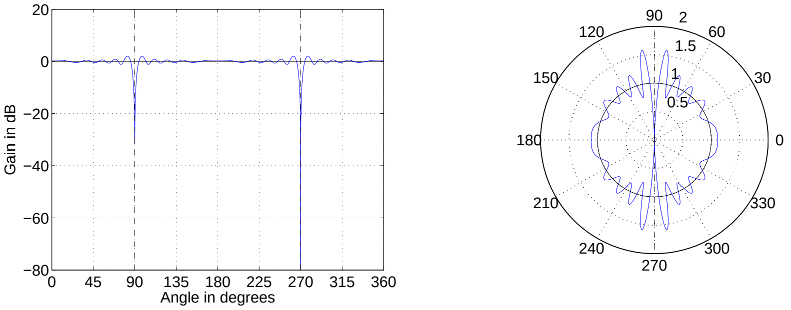

(e) Attenuation/amplification in dB. (f) Linear antenna gain.

Fig. 2. Sample antenna gain pattern of a ULA10 with varying directions (40°, 80°, 90°) of one antenna null. Other nulls (at 320°, 280°, and 270°) result from the symmetry of linear array antennas.

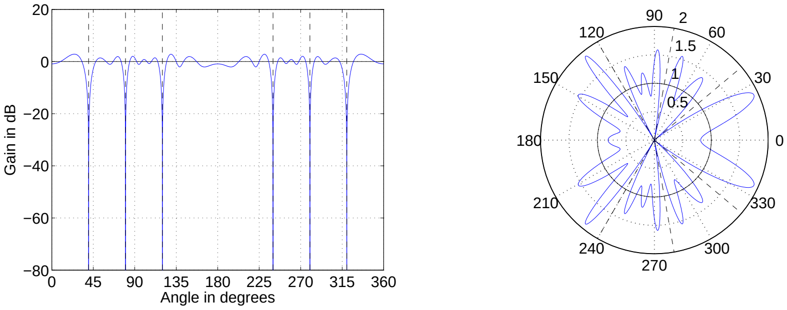

(a) Attenuation/amplification in dB. (b) Linear antenna gain.

Fig. 3. Antenna gain pattern of a ULA10 with three antenna nulls in the directions 40°, 80°, 120°.

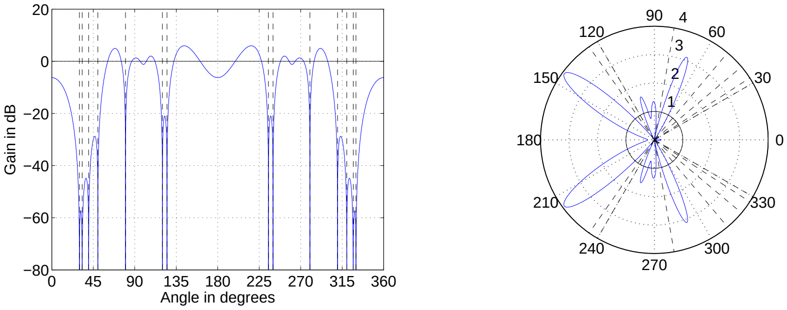

(a) Attenuation/amplification in dB. (b) Linear antenna gain.

Fig. 4. Antenna gain pattern of a ULA10 with seven antenna nulls in the directions 30°, 33°, 40°, 50°, 80°, 120°, 125°.

4 Analysis of Interference Reduction

This section provides an analysis of the above-described approach by means of simulations. We analyze the interference cancellation capabilities of adaptive antennas in connection with the proposed nulling algorithm, and the resulting benefits in terms of spatial reuse.

4.1 Neighborhood Relationships and Null Placement

With our nulling approach, not all contiguous interference sources can be nulled in arbitrary traffic relationship constellations. Possible gains in spatial reuse are thus not only depending on the nulling capabilities of the adaptive antennas, but also on topology and traffic scenarios.

When placing antenna gain nulls based on neighborhood communication patterns several questions arise: How many desired and undesired neighbors does a node typically see? Does the presence of desired nodes inhibit null placement, in particular since the limited number of antenna elements of the antenna array do not allow for arbitrary antenna gain shaping (cf. Fig. 4)? Can the proposed nulling approach in connection with the chosen beamforming method effectively suppress interference?

To give an answer to these questions, we performed extensive simulations with the models as described in Sec. 3. In the random network with 50 nodes, we show results for both 10 concurrent data flows and 50 concurrent data flows. In the latter case all 50 nodes set up a multi-hop data flow to randomly chosen sinks, resulting in a scenario with highly meshed traffic relationships. In particular in such a case we must ask whether antenna nulls can be placed at all while preserving all links to desired neighbors.

Let us begin with an analysis of how many neighbors are desired nodes, and how many neighbors a node wishes to null. As described in more detail in Sec. 2, a neighbor of a given node is called a *desired* if the node either transmits to or receives from that neighbor in connection with any of the data flows in the network. All neighbors that are not desired nodes are called *undesired*. The probability mass functions of the number of desired and undesired nodes for a node in the network are shown in Fig. 5.

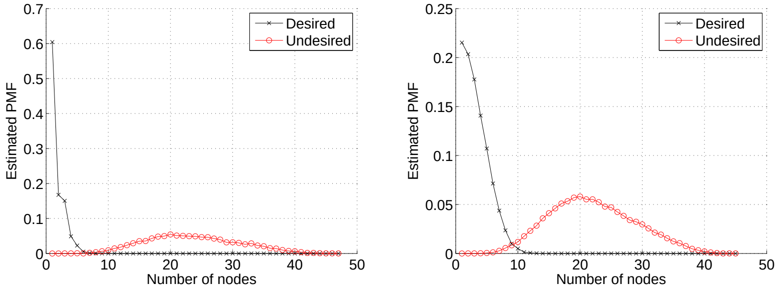

(a) 10 data flows. (b) 50 data flows.

Fig. 5. Probability mass function of the number of desired and undesired nodes.

It is now interesting the analyze how many nulls a node actually forms under the constraints of preserving desired links and limited degrees of freedom of the antenna array. The probability mass function of the number of placed nulls is shown in Fig. 6. Interestingly, even in the case of 50 active data flows, many nodes fully use the degrees of freedom (= 9) of the ULA10 for nulling.

4.2 Node Decoupling and Interference Suppression

Each antenna gain null is placed explicitly toward an undesired node. However, we expect that each placed null will on average suppress more than one

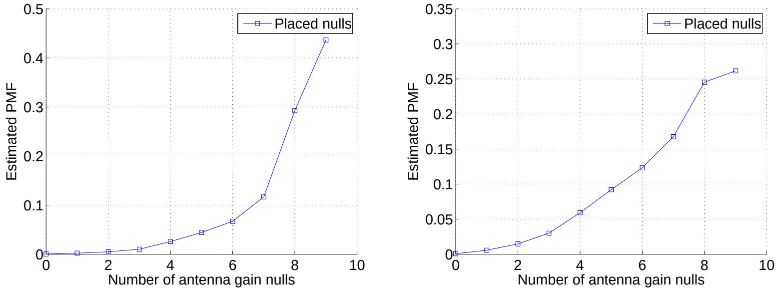

(a) 10 data flows. (b) 50 data flows.

Fig. 6. Probability mass function of the number of placed nulls.

interferer lying in the angular direction of the null. Fig. 7 shows the probability mass function of the number of nodes in sensing range before and after nulling.

As can be seen, this number is reduced dramatically. From these results we expect a significant reduction in MAC layer blocking, by explicit control messages or by carrier sensing, and a corresponding increase in the spatial reuse of radio resources.

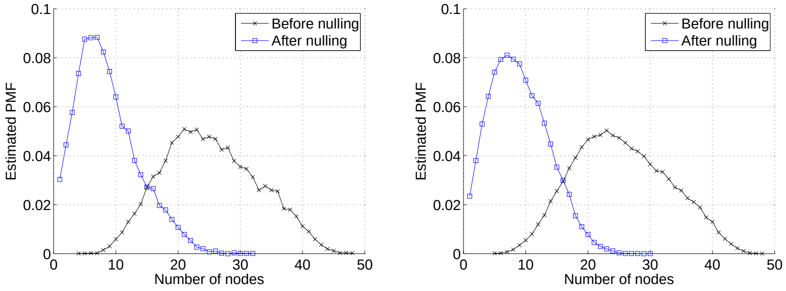

(a) 10 data flows. (b) 50 data flows.

Fig. 7. Probability mass function of the number of nodes in sensing range.

We are also interested in the interference reduction, change in signal level of the desired signal, and SIR improvement provided by nulling. As can be seen in Fig. 8 and Fig. 9, nulling both decreases interference and increases the signal level of the desired signal, on average. Here, the notion of interference is the worst-case interference level, which occurs when all remaining nodes which are part of at least one of the multi-hop paths transmit simultaneously. As a result of increased signal level and reduced interference, the worst-case SIR (Fig. 10) benefits significantly from nulling.

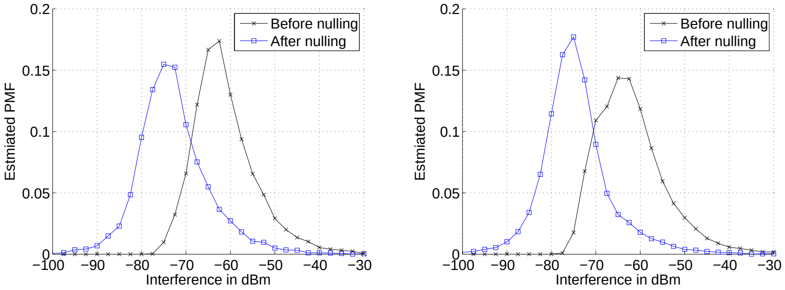

(a) 10 data flows.
(b) 50 data flows.

Fig. 8. Probability mass function of the worst-case interference level.

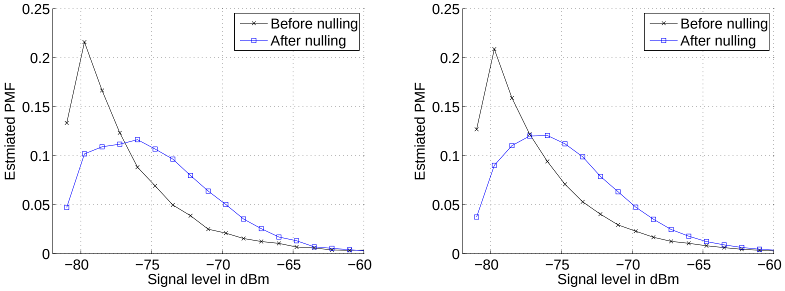

(a) 10 data flows.
(b) 50 data flows.

Fig. 9. Probability mass function of the signal level of the desired signal.

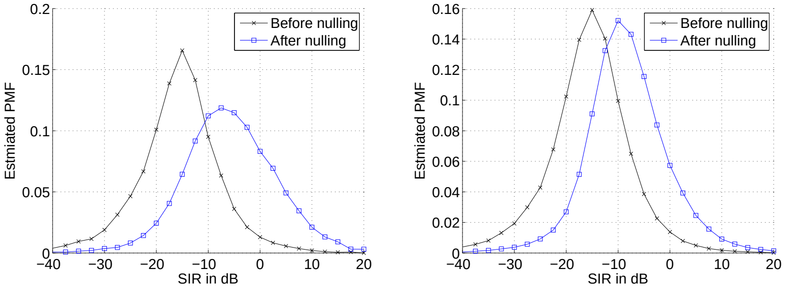

(a) 10 data flows.
(b) 50 data flows.

Fig. 10. Probability mass function of the worst-case signal-to-noise ratio.

5 Aspects of Connectivity and Routing

The previous sections showed how placing antenna gain nulls can drastically reduce interference. Thereby, our nulling approach tries to ensure that desired links persist. Antenna nulls may still prohibit desired links in case of movement or rotation of mobile nodes, changes in communication relationships, and other reasons for necessary changes of communication paths, such as shadow fading. The involved loss of desired links might induce losses in terms of connectivity, and result in an increase of path lengths. Means to counteract these shortcomings could be as follows.

- Updating antenna nulls by repeating the process of neighborhood exploration and DoA estimation. Factors mandating such a re-computation of null directions have been summarized in Sec. 2.
- Regular or tentative removal of some or all antenna nulls, as a conservative means to avoid long-term disconnection and routing detours.
- Analysis of routing protocol control packets: With certain routing protocols, a node may infer that a (control) packet has traversed one of its nulled neighbors. Removing the corresponding antenna null can eliminate the obviously existing route detour.

In the following, we will go one step back by asking the following question: With our nulling approach, do connectivity and path lengths deteriorate at all for the network scenario considered in this work? If so, how much?

We answer this question by dynamically changing traffic relationships in the network. In particular, upon adapting antenna nulls to desired and undesired nodes in the neighborhood, we let each node *reselect* its chosen sink. While the prevailing nulls have been placed in response to the initially chosen sinks and the resulting communication paths, the new sinks now have to be reached via the links that remained after this null placement.

Upon reselecting sinks and performing routing to these sinks, we will numerically analyze the resulting path length distribution. We then compare the results to the path length distribution we would have obtained if no nulls would have been placed at all.

In the same network setup as considered above, all 50 nodes set up a data flow toward a randomly chosen sink, separately. The results for the path length distribution for the reselected sinks are shown in Fig. 11. On the horizontal axis, a number of hops equal to zero means that the sink cannot be reached (disconnected), and a number of hops equal to one means that the sink is a one-hop neighbor. The curve labeled “Without sinks” represents the path length distribution that would be obtained if the antenna nulls placed upon selecting the initial sinks would be removed before routing to the reselected sinks.

Surprisingly, it can be observed that nulling does not deteriorate the path length distribution. Less long paths exist when the nulls based on the initially chosen sinks/routes are maintained. This means that although the antenna

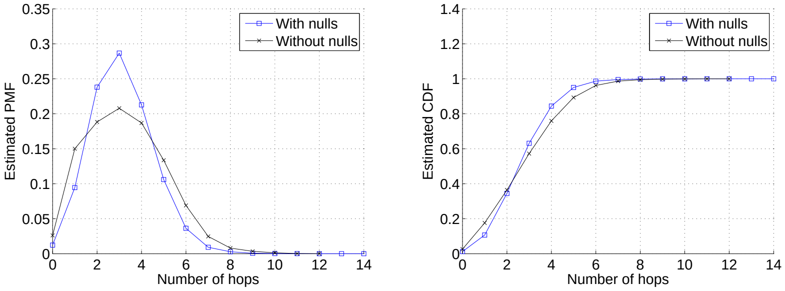

(a) Estimated probability mass function.

(b) Estimated cumulative distribution function.

Fig. 11. Path length distribution from source to reselected sink.

nulls have been placed for different traffic relationships, the same antenna gain patterns let the new sinks often be reached in a smaller number of hops.

This result can be explained by looking at the antenna patterns depicted in Sec. 3, and relating to previous work on so called *Random Direction Beamforming* [2] in multi-hop networks. Antenna gains exceeding a factor of one can lead to very long links in the network. These long links can provide shortcuts toward the sink which are not available with omni-directional antennas. On the other hand, the number of neighbors (node degree) typically decreases both with gain maximizing beamforming and interference nulling. For this reason, the number of one hop paths is shorter with nulling than with omni-directional patterns (cf. Fig. 11, data point for one hop).

The means to counteract disconnection listed above may still be necessary in practice. Nevertheless, the results shown here indicate that for many scenarios the proposed nulling scheme does not inhibit a well-connected multi-hop network, even when nulls are not up-to-date with respect to current communication paths.

6 Conclusions and Future Work

In this paper we proposed algorithms to adapt the antenna patterns of nodes with smart antennas based on the long-term communication pattern observed in a node's neighborhood. This provides a useful trade-off between complexity and performance, and may be more useful in practice than per-packet beamforming in connection with per-packet Direction of Arrival (DoA) estimation.

The provided simulation results indicate that the proposed scheme is beneficial in terms of interference and the decoupling of nodes contending for radio resources. In case of outdated neighborhood information, spuriously placed antenna gain nulls do not affect the connectivity of the wireless network.

For future work, we intend to extend our algorithms to also support nulling of interference from neighbors outside the communication range, and external interference from other systems. Further, we are interested in topological effects of using nulling on top of gain maximizing beamforming as used in previous work. In our current work, we go beyond topological analysis and evaluate the MAC layer performance when using the proposed nulling scheme. Initial results show that the benefits described here also translate into remarkable improvements in terms of the end-to-end throughput.

Acknowledgments

The authors would like to thank Diego Herrero for implementing the adaptive beamforming algorithms used in this work.

References

1. R. Choudhury, X. Yang, R. Ramanathan, and N. Vaidya, "On designing MAC protocols for wireless networks using directional antennas," *IEEE Transactions on Mobile Computing*, vol. 5, May 2006.
2. R. Vilzmann, C. Bettstetter, D. Medina, and C. Hartmann, "Hop distances and flooding in wireless multihop networks with randomized beamforming," in *Proc. ACM Intern. Workshop on Modeling, Analysis, and Simulation of Wireless and Mobile Systems (MSWiM)*, (Montreal, Canada), Oct. 2005.
3. N. Malhotra, M. Krasniewski, C. Yang, S. Bagchi, and W. Chappell, "Location estimation in ad hoc networks with directional antennas," in *IEEE International Conference on Distributed Computing Systems (ICDCS)*, June 2005.
4. R. Vilzmann, J. Widmer, I. Aad, and C. Hartmann, "Low-complexity beamforming techniques for wireless multihop networks," in *Proc. of IEEE Commun. Soc. Conf. on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, Sept. 2006.
5. M. Takai, J. Martin, A. Ren, and R. Bagrodia, "Directional virtual carrier sensing for directional antennas in mobile ad hoc networks," in *Proceedings of MobiHoc*, 2002.
6. T. Korakis, G. Jakllari, and L. Tassiulas, "A MAC protocol for full exploitation of directional antennas in ad-hoc wireless networks," in *Proceedings of MobiHoc*, 2003.
7. S. Roy, D. Saha, S. Bandyopadhyay, T. Ueda, and S. Tanaka, "A network-aware MAC and routing protocol for effective load balancing in ad hoc wireless networks with directional antenna," in *Proceedings of MobiHoc*, 2003.
8. Z. Shan and T.-S. Yum, "A conjugate augmented approach to direction-of-arrival estimation," *IEEE Transactions on Signal Processing*, vol. 53, pp. 4104–4109, Nov. 2005.
9. H. Steyskal, "Synthesis of antenna patterns with prescribed nulls," *IEEE Communications Magazine*, no. 2, pp. 273–279, 1982.