



<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *IEEE International Workshop on Factory Communication Systems, Lemgo/Detmold, Germany, 21-24 May, 2012.*

Citation for the original published paper:

Willig, A., Uhlemann, E. (2012)

On relaying for wireless industrial communications: Is careful placement of relayers strictly necessary?.

In: *IEEE International Workshop on Factory Communication Systems - Proceedings, WFCS* (pp. 191-200). Piscataway, NJ: IEEE Press

<http://dx.doi.org/10.1109/WFCS.2012.6242567>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:hh:diva-19678>

On Relaying for Wireless Industrial Communications: Is Careful Placement of Relays Strictly Necessary?

Andreas Willig
University of Canterbury
Dept. of Computer Science and Software Engineering
Christchurch, New Zealand
andreas.willig@canterbury.ac.nz

Elisabeth Uhlemann*
Halmstad University
CERES: Center for Research on Embedded Systems
Halmstad, Sweden
elisabeth.uhlemann@hh.se

Abstract

*Relaying is a very promising technique to improve the reliability of data transmission in wireless (industrial) networks. With relaying, relay nodes support source nodes in carrying out retransmissions. A common assumption is that relayers should be placed at “good” positions (e.g. in the middle between source and destination) to achieve benefits. In this paper we tackle the question of whether it is strictly necessary to place relayers at “good” positions (which often requires extensive measurements). We present results indicating that the benefits of relaying are achievable even with randomly placed relayers, as long as enough of them are deployed. Specifically, we present results suggesting that with a sufficient (and still not too high) number of randomly deployed relayers, the probability that **all** packets, sent by source nodes to a central controller in a TDMA round, reach the controller is larger than for the case with source-only retransmissions. This finding holds true both in the absence and the presence of feedback.*

1. Introduction

A key challenge in the design of wireless industrial communication systems is to transmit packets reliably despite the presence of channel noise [1]. To cope with deadlines, many industrial communication systems, including the recently standardized Wireless HART technology [2], [3] adopt a TDMA-based scheme on the MAC layer. The handling of transmission errors often involves packet retransmissions, for which in a TDMA-based setting extra

time slots have to be reserved. A standard assumption is frequently that any retransmissions are carried out by the original sender itself, so that one and the same source node gets allocated a number of slots to use for its original transmission and any retransmissions or packet repetitions.

In the last few years the communications and networking research community has looked intensively at so-called *cooperative communication schemes*, which rely on exploiting the spatial diversity of the wireless channel [4], [5], [6], [7], [8], [9]. These schemes suggest a new approach in the design and implementation of retransmission-based error control protocols: instead of letting the source of a data packet perform all retransmissions, it often makes sense to involve third-party nodes in the process. For example, a data packet sent by a source node might have been received erroneously by the intended receiver, but it might have been overheard (i.e. received successfully) by a third node, the **helper node** or **relay node**. This is possible because the wireless transmission medium is a broadcast medium. Then, the relay node might perform a retransmission on behalf of the source, when it realizes that the destination has not received the source packet from the source. Or, when the overall protocol does not use any feedback, the relayer can perform a repeated packet transmission instead of the source.

It is intuitively clear (and substantiated by several publications, including [10]) that such an approach is especially promising if the relay has “good” channels to both source and destination, like for example when a relay is placed in the middle between the source and the destination. However, in many industrial settings it might for practical reasons not be possible or easy to place relay nodes at a desired position, for example because this position is inaccessible or dangerous to human staff. Further-

*E. Uhlemann is partly funded by the Swedish Governmental Agency for Innovation Systems, Vinnova, through the VINNMER program, www.vinnova.se

more, what is a “good” position depends on the characteristics of the wireless channels between the involved nodes, which need to be measured carefully. Such measurements can require a substantial amount of time and money.

In this paper we consider uplink traffic in a simple, TDMA-based wireless industrial network in which a number of sources, supported by a number of relayers, wish to transmit data to a central controller. In our considered TDMA scheme, time is sub-divided into subsequent superframes, which in turn are sub-divided into a fixed number of time slots. Each source wishes to transmit one packet during each superframe. There are sufficient time-slots in a superframe and hence also before the deadline to allow for initial transmissions from a source to the controller, plus a number of redundant transmissions (made either by the source nodes or by relay nodes, who may have overheard previous transmissions of the source or other relayers) with the aim to improve the chances that the source data is indeed received by the controller.

The key point of our study is that we do not make any assumptions on “good” placement of relay nodes. Instead, we want to know whether there are still advantages to relaying when both the sources and the relays are placed completely **randomly**. We present simulation results indicating that already a relatively small number of randomly placed relays carrying out all re-transmissions on behalf of the sources is sufficient to substantially improve the probability that **all** source packets in a superframe reach the central controller. This finding is true both in the complete absence of feedback from the central controller, and in the presence of perfect feedback. While we do not argue against the value of a “good” deployment, this can in practical scenarios sometimes be hard to achieve. Our results show that even with randomly deployed relayers, the relaying approach still provides performance benefits, provided a sufficiently large number of relayers is deployed.

This paper is structured as follows: In the next Section 2 we provide a brief background on spatial diversity and cooperative communications concepts. In Section 3 we present our system model. Following this, in the next two sections we provide the main results of this paper. Namely, in Section 4 we provide simulation results for the case without feedback, whereas the case with perfect feedback is covered in Section 5. Related work is discussed in Section 6 and our conclusions are given in Section 7.

2. Background: Spatial Diversity and Cooperative Communications

Generally, wireless data transmissions can suffer from phenomena like path loss and shadowing, multi-path propagation and thermal noise [11], [12], [13]. These propagation phenomena lead to degraded signal strength and signal quality at the receiver. When additionally there is mobility of the nodes or in the surrounding environment, further signal degradation can be incurred by time-varying self-interference and (fast) fading phenomena.

A key approach to improve transmission quality are diversity schemes. Generally speaking, diversity schemes aim to provide the receiver with multiple, ideally independently faded copies of the same signal, which the receiver can suitably combine. There are several distinct diversity schemes, which can be broadly classified into time diversity, frequency diversity and spatial diversity. In spatial diversity schemes [4] information is transmitted over multiple geographically separated antennas. In modern multiple-input, multiple-output (MIMO) systems like IEEE 802.11n the transmitter and receiver both have multiple antennas, whereas in cooperative communication schemes the additional antennas are provided by third-party nodes [5], [6], [7], [8]. In cooperative schemes all involved nodes can be single-antenna nodes. Following the terminology developed for one particular cooperative scenario, the relay channel, we refer to these third-party nodes as **relayers**.

The general concept of cooperative communications then gave rise to the idea of incorporating relay nodes into retransmission-based error-control schemes – see the incremental relaying scheme proposed in [7], and see [14] for an example in an industrial networking context. In this kind of schemes a relayer can perform a retransmission on behalf of a source node, provided that the relay node has overheard the original data packet from the source.

3. System Model

3.1. Network Setup and Channel Model

We consider a wireless industrial communication system with $N = M + K + 1$ nodes as follows:

- M source nodes, numbered from 0 to $M - 1$
- K relay nodes, numbered from M to $M + K - 1$
- One central controller, denoted $M + K$.

In this work we consider only uplink traffic, i.e. traffic from the source nodes to the central controller. This is a typical setup in wireless industrial communication systems, in which sensors report their readings to a central controller. The basic goal is to transmit data packets reliably from the sources to the controller, perhaps with the help of the relayers. For simplicity we assume that source nodes and relay nodes are truly distinct nodes. We further assume a TDMA setup. Time is subdivided into superframes, which in turn are sub-divided into $T \geq M$ time slots. One time slot is large enough to accommodate one packet, and, in the presence of feedback from the central controller, also a suitable feedback signal. All source packets have the same size. The source packets include at minimum the source identifier, a packet checksum (for example, a CRC code) and source data. For simplicity, we do not consider any kind of coding, network coding or packet combining, but the CRC is assumed to be perfect,

so that all errors are detected reliably. When the destination or a relay receives a packet with a wrong CRC value, it discards the packet.

Each source node generates a new data packet at the beginning of each superframe. In the first M of the T slots each source transmits its packet in turn. The relayers listen in all T slots and try to overhear the packets sent by the sources or the other relayers. As a result of this, the relayers can accumulate more and more information in the course of a superframe. The remaining $T - M$ slots are used for retransmissions. More precisely, each of these $T - M$ slots is allocated exclusively to either a source node or a relay node, and this allocation does not change over time. If a slot is allocated to a source node, the source simply retransmits its own packet. If it is allocated to a relay, the relay picks one of the source packets it has overheard so far (which is a random subset of all source packets) and retransmits this packet. Under these assumptions our setup amounts to the use of (possibly distributed) repetition coding.

We use a conceptually simple time-varying channel model. We assume a shared wireless medium. Between each pair of nodes there exists a separate, symmetric wireless channel that is stochastically independent of all other channels. Each channel is characterized by its packet loss rate, which is defined with respect to the packet size used by the source nodes. These packet loss rates are constant throughout a superframe, but change between superframes. The channel packet loss rates at the beginning of superframe t can be summarized in an $N \times N$ symmetric matrix $\mathbf{C}_t = ((c_{i,j}(t)))_{i,j \in \{1, \dots, N\}}$, where $c_{i,j}(t)$ denotes the packet loss rate on the symmetric channel between nodes i and j throughout superframe t . At the beginning of each new superframe t , a new random realization of the symmetric channel matrix \mathbf{C}_t is generated, where each entry $c_{i,j}(t)$ is taken from a uniform distribution between 0 and 1. With these assumptions, our channel model can be regarded as a block-fading channel. None of the transmission schemes considered in this paper assumes any kind of channel knowledge.

3.2. Performance Metric

The main performance metric considered in this paper is the probability that at the end of a superframe the central controller possesses the packets of **all** sources. We refer to this as the **success probability**. When at least one source packet is missing, we refer to this as a **failure**.

4. Results in the Absence of Feedback

In this section we present results for a system setup without any feedback from the central controller. The main control knob that we have considered is how the relayers actually pick the packets they retransmit from the (random) subset of source packets they possess at the beginning of their given timeslot. We specifically consider schemes where the relayers make this choice randomly:

- In the **uniform scheme** each packet is chosen with the same probability.
- For the **inverse scheme** each relay keeps for each source packet a counter for how often it has already successfully received this packet from either the source or another relay. These counters are initialized to zero at the beginning of a superframe. The intuition behind the inverse scheme is that a relay favors packets which it has not heard so often over packets which it has heard more often (and which hence possibly have a higher chance of already having been received by the controller). When at the beginning of its slot the relay does not possess any packet, it keeps quiet. Otherwise, when the relay has overheard source packet i already $c_i \geq 0$ times, it assigns to this packet the probability

$$p_i = \frac{c_i}{\sum_{j=1}^M c_j}$$

and the actual packet to be retransmitted by the relay is chosen according to this probability distribution.

Both schemes allow for a very simple system setup: there is no need for feedback from the controller, no node needs to know channel characteristics and the sources do not need to be aware of the presence of relayers.

In the following we have for a given number M of sources constructed schedules in which each source has an additional R slots for retransmissions, so that a superframe has $T = (R + 1) \cdot M$ slots in total. When there are no relayers (i.e. $K = 0$), the schedule of the superframe consists of $R + 1$ repetitions of $\{0, 1, \dots, M - 1\}$, i.e. for example with $M = 3$ sources and $R = 3$ retransmissions the schedule becomes

$$0, 1, 2, \quad 0, 1, 2, \quad 0, 1, 2, \quad 0, 1, 2$$

When there is at least one relay, the first M slots of the schedule are allocated to the sources, whereas the last $T - M$ slots are allocated to all relayers in a cyclic fashion. For example, with $M = 3$ sources, $K = 4$ relayers and $R = 5$ retransmissions the schedule becomes

$$0, 1, 2, \quad 3, 4, 5, \quad 6, 3, 4, \quad 5, 6, 3, \quad 4, 5, 6, \quad 3, 4, 5$$

which has a total of $M \cdot (R + 1) = 18$ slots.

We have obtained the success probability using a specifically tailored simulator. For each setting of simulation parameters (M , K , R , uniform vs. inverse scheme) we have simulated 30,000 superframes so that for each new superframe a completely new and independent channel matrix is generated. Within each such superframe the simulator starts first by letting all the sources transmit a packet. With the help of the channel matrix it is then determined which relayers (and the controller) have received which source packets. Subsequently, the relay transmissions are handled, where again the channel matrix is consulted to determine if the other relayers (or the destination)

have picked up the transmission of the currently transmitting relay. At the end of a round it is checked whether the central controller has received all frames. Since the response variable (success or not) of a single superframe can be modeled as a Bernoulli random variable, the confidence interval half-widths for a confidence level of 95% can according to [15, p. 417] be upper-bounded by ≈ 0.0057 [15, p. 417] and are not shown in the figures.

In Figure 1 we show the success probability for the case of two sources ($M = 2$), a varying number of retransmissions R and a varying number of relayers K . We show similar results for the case of six sources in Figure 2 and for twelve sources in Figure 3. The following points are noteworthy:

- In all cases, the average success probability of the inverse scheme is significantly better than the success probability obtained with the uniform scheme. This is likely caused by the ability of the inverse scheme to focus the effort of relayers with later slots to those source packets which have not been retransmitted so often and hence are less likely to have been received successfully by the central controller.
- In all considered cases, when there are at least four relayers present and at least two retransmissions can be carried out, it is better to let relayers do all the re-transmissions instead of the source nodes. Hence, four relayers can be considered as a “critical mass”, giving a reasonably high chance that there is at least one relay having a good-quality channel to the central controller. As a result of information accumulation, this relay in turn has good chances to pick up packets from sources and other relayers.
- When fixing the number of relayers to $K = 4$, the improvement in success probability for two sources is only minor, but as the number of sources increases, the advantage of using four relayers instead of using the sources becomes more pronounced.
- The performance difference between the cases of four and eight relayers are significant for both the inverse and the uniform scheme, whereas the difference between eight and ten relayers is much smaller. This suggests that there is a saturation point for the number of relayers.
- When eight or ten relayers and the inverse scheme are used, the success probability comes close to one much more quickly than in the case without relayers with increasing numbers of retransmissions.

Limited as they are, these results clearly demonstrate that, given a sufficient number of relayers (at least four), it is better to let the relayers carry out all retransmissions than to leave this to the sources. In a practical setup, this allows to remove complexity and energy expenditure (for retransmissions) from source nodes, at the cost of having

to deploy additional relay nodes. Furthermore, our results indicate that not much care is needed during the deployment process to achieve a gain with relayers, allowing it to be less cost-intensive than it would be when careful placement is exercised.

5. Results with Feedback Present

In this section we look at a similar scenario as before, but now we consider the presence of perfect feedback from the central controller. While this is certainly an idealizing assumption, it is generally not impossible to come close to having perfect feedback, for example by assuming that the central controller uses a high transmit power for its feedback signal.

We use the feedback in different ways, depending on whether we use source-retransmissions or employ relayers. When relayers are employed, we start by letting each source transmit its packet once. After finishing this first round of source transmissions, the central controller issues the first feedback packet, indicating which source packets it has received. When receiving this, the relayers drop the packets that the controller indicated, so that they will not be picked anymore by a relay in any of its slots. Following this, the relayers transmit according to the schedule described in Section 4. After each relay transmission the controller again sends a feedback packet as before, which is processed by the relayers. Please note carefully that this scheme allows more efficient use of the retransmission slots: when the central controller receives a source packet quickly (for example directly from the source), no retransmission slots are spent for this source anymore, and they can be used for the packets of other sources.

When the sources are doing all retransmissions themselves, the feedback operation works slightly different. We have adopted the queued-retransmission scheme described in [16], which achieves similar re-use of slots as in the case with relayers. First, all sources transmit once. At the end of these M transmissions, the controller checks which source packets it has and puts the source addresses of missing packets into a list. At the beginning of each overhead slot (these are the $R \cdot M$ slots allocated for retransmissions) the controller takes the head of the missing list, say source i , and informs source i that it is now its turn (we assume that the packets by which the controller selects the next source do not take additional bandwidth or time). Source i transmits and the controller checks whether the packet has been correctly received. If not, i remains in the missing list and the operation continues. Otherwise, i is dropped and the operation goes on with the remaining items in the missing list.

In Figure 4 we show the success probability for the case of two sources ($M = 2$), varying number of retransmissions and varying number of relayers. We have included both results with and without feedback for comparison purposes. For the relay-based schemes, we have only considered the inverse scheme (see Section 4). In Figure 5 and

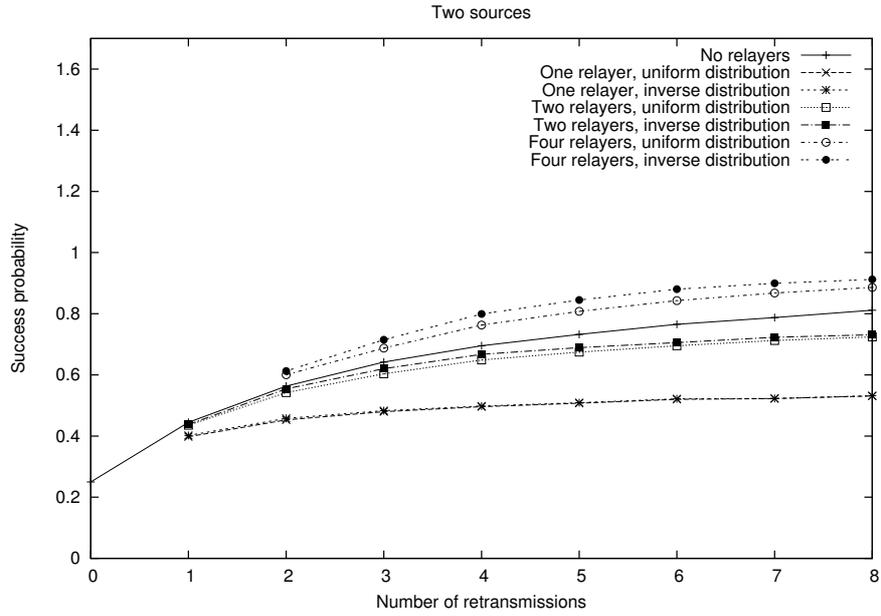


Figure 1. Average success probability with $M = 2$ sources, varying numbers K of relayers, varying numbers R of retransmissions and no feedback.

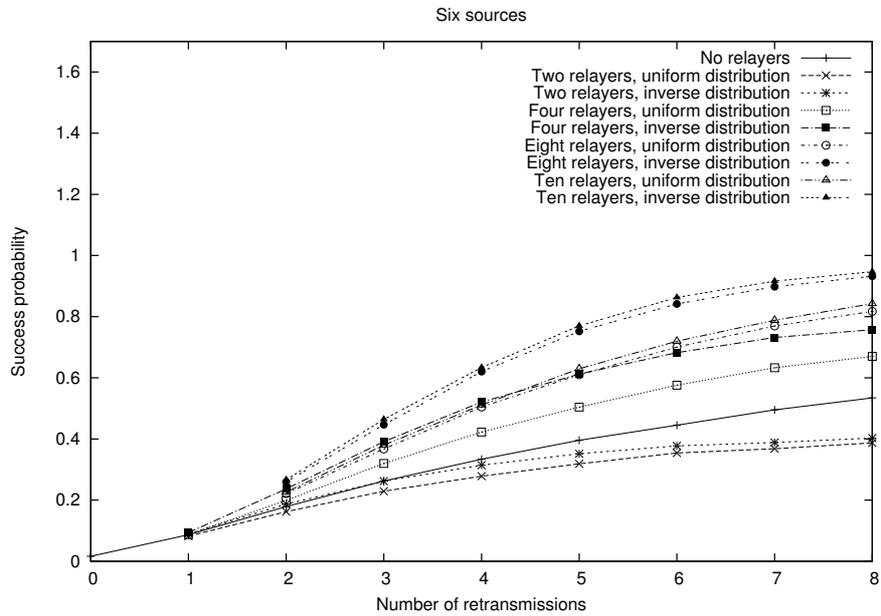


Figure 2. Average success probability with $M = 6$ sources, varying numbers K of relayers, varying numbers R of retransmissions and no feedback.

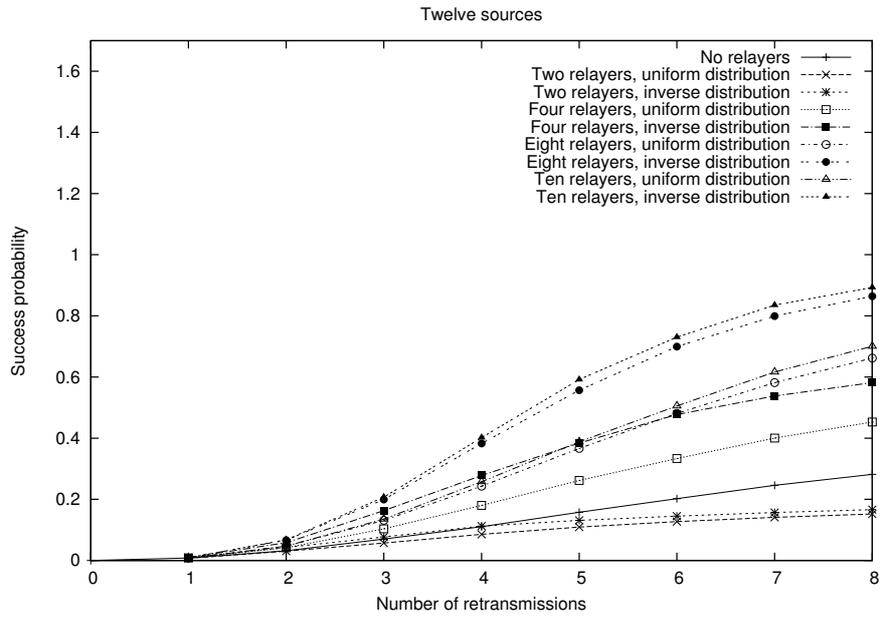


Figure 3. Average success probability with $M = 12$ sources, varying numbers K of relayers, varying numbers R of retransmissions and no feedback.

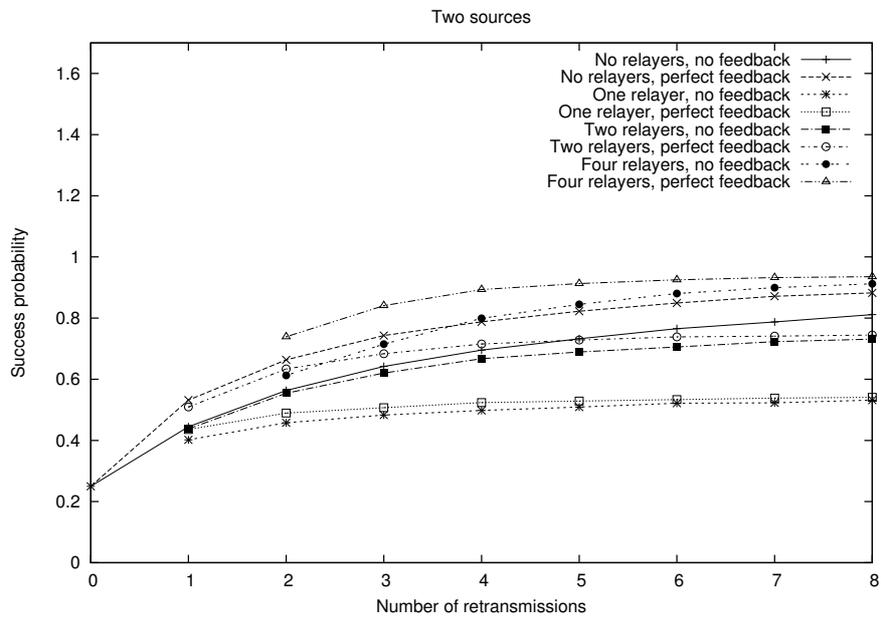


Figure 4. Average success probability with $M = 2$ sources, varying numbers K of relayers, varying numbers R of retransmissions, perfect feedback and using the inverse scheme.

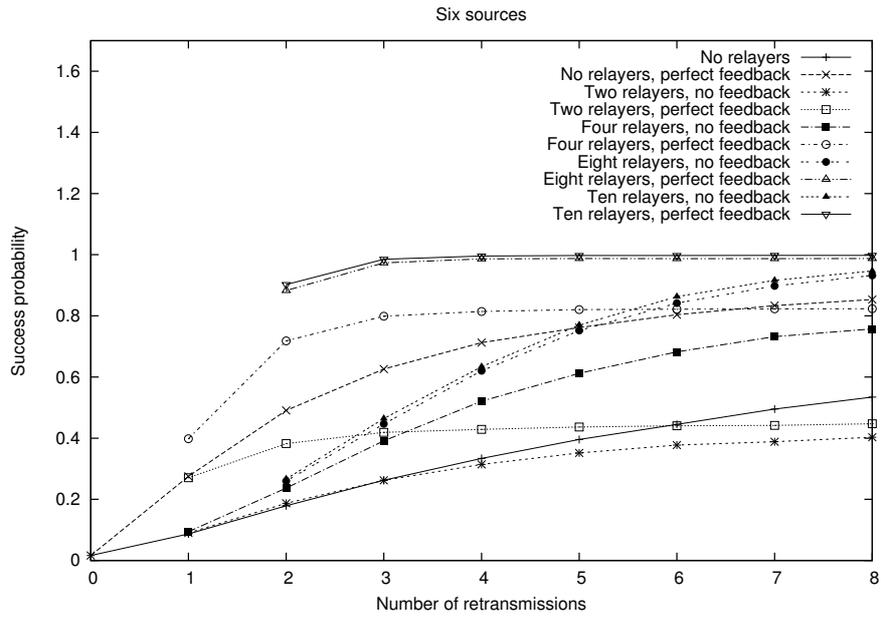


Figure 5. Average success probability with $M = 6$ sources, varying numbers K of relayers, varying numbers R of retransmissions, perfect feedback and using the inverse scheme.

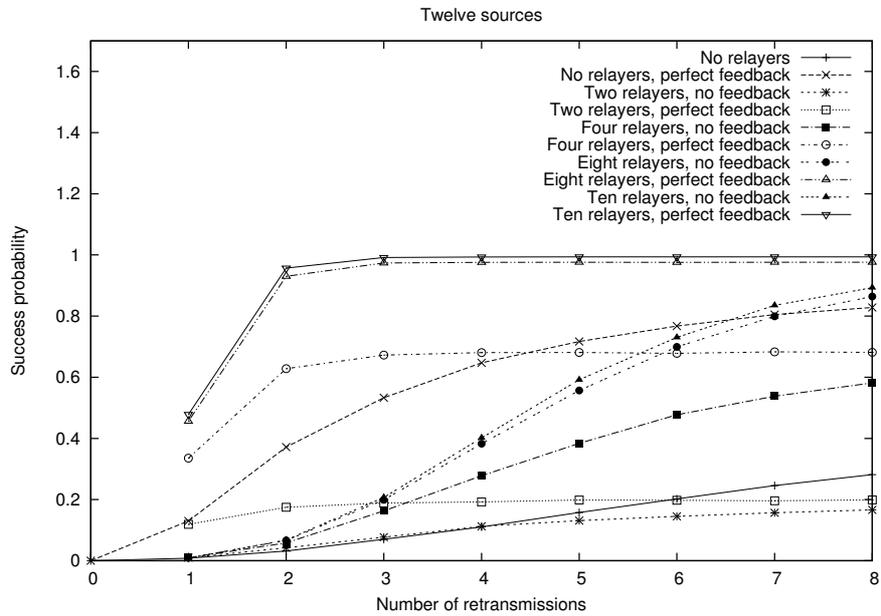


Figure 6. Average success probability with $M = 12$ sources, varying numbers K of relayers, varying numbers R of retransmissions, perfect feedback and using the inverse scheme.

Figure 6 we show similar results for the cases of six and twelve sources, respectively. The following findings are interesting:

- For all considered numbers of sources there is a number of relayers beyond which it is better to use relayers than to rely on the sources. For two sources, four relayers are needed, for six and twelve sources eight relayers are needed to consistently outperform the scheme with only source retransmissions.
- With eight or ten relayers and at least three retransmissions ($R = 3$) it becomes possible to reach a success probability close to one. In contrast, without relayers the success probability converges, as expected, to one with increasing number of retransmissions, but is still substantially away from one within the considered range from one to eight retransmissions.
- For both six and twelve sources, the difference between eight and ten relayers is very small, whereas the difference between four and eight relayers is significant.
- When relaying is used, for all considered numbers of relayers the success probability appears to converge to a fixed value with increasing number of retransmissions – when feedback is present the limiting value is reached much more quickly, though. In other words, adding retransmissions does not help beyond a certain point. A likely explanation for this is that the overall success probability is clearly bounded by the ability of the set of relayer to capture *all* source packets – with too few relayers this probability of full capture is below one and cannot be improved by adding retransmissions.
- For a fixed number of relayers and a fixed number of retransmissions, the scheme with perfect feedback is (expectedly) significantly better than the scheme without feedback. This gives an insight into the value of feedback.

From these findings we conclude that the presence of feedback does not substantially alter our findings from the case without feedback: there is a threshold number of relayers (eight in the cases considered here) beyond which it is better to use relayers instead of source retransmissions.

We finally look at our results from a different perspective. In Figure 7 we show results for $M = 20$ sources, a fixed number $R = 6$ of retransmissions and varying numbers of relayers. We consider both the cases without feedback and with perfect feedback. Furthermore, the figure shows (as straight lines) the success probability for systems without relayers, both without feedback and with perfect feedback. The following results are notable:

- If we consider the case without any feedback, already three relayers are sufficient to achieve a better probability that all $M = 20$ source packets reach the

central controller than can be achieved with retransmissions by the sources themselves. Adding further relayers leads to further improvements, but the success probability appears to converge to a value significantly smaller than one.

- When feedback is present and six or more relayers are used, relaying has an advantage over the case without relaying. However, once eight or nine relayers are used (less than half of the number of sources!) the success probability is very close to one.

6. Related Work

Relaying in general is a well-established research topic, and many results, both theoretical and practical have been published to date, see, e.g. [5] and [6]. It is well-known that multi-antenna techniques can, due to spatial diversity, improve the achievable rates or the transmission reliability. In relaying and other cooperative techniques, the multiple antennas are distributed over several involved nodes, thus creating the need for coordination of activities of sources and relayers. The benefits of relaying for error-control purposes in industrial environments have for example been discussed in [1] in general. In [14] a framework is presented in which relayers compete with each other to support one single source node. By measuring their own success (both in winning relay slots against other relayers and in successfully transmitting packets to the receiver) the relayers can adapt their willingness to help and over time the best (few) relayers for a fixed source emerge. In the approach presented in [17] it is a source node who selects its relayers based on a systematic trial-and-error procedure.

Some of the important questions in the design of relaying schemes include the placement of relay nodes [10] and the allocation of relayers to sources, see e.g. [18], [19], [20]. However, many of these publications focus either on network-wide throughput, the probability of error for individual nodes, or the average error probability taken over all nodes (see for example [10], which use a much more elaborated channel model than the one used in this paper). In an industrial context, even when restricted to the uplink case, it is however more important to successfully receive the data from *all* sources, and often this is even coupled with delay constraints (deadlines). This naturally leads to the choice of performance measures different from the average error probability (taken over all nodes).

7. Conclusions

In this paper we have presented simulation results suggesting that for periodic TDMA uplink traffic, it might be more beneficial to use relayers than to depend on source retransmissions, provided that there are a sufficient number of relayers and also a sufficient number of slots for retransmissions. While this finding as such is not so sur-

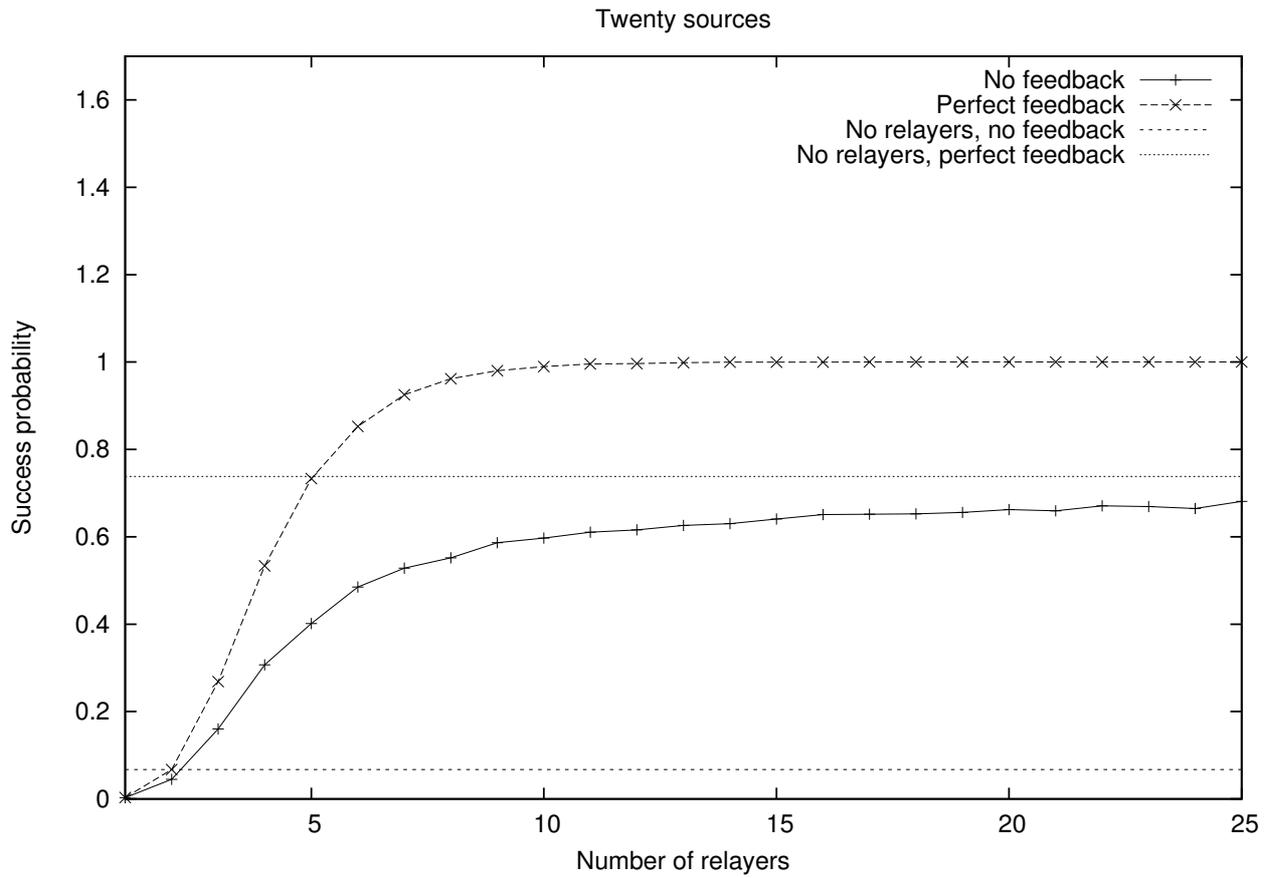


Figure 7. Average success probability with $M = 20$ sources, varying numbers K of relayers, a fixed number of $R = 6$ retransmissions, and with both no and perfect feedback (In the absence of feedback the relayers use the inverse scheme).

prising in itself, the fact that this is true even for **completely random** placements of (not excessively many) relayers was not obvious before.

There are some opportunities for further research. A first shortcoming of the present paper is that our choice of uniformly distributed packet loss rates in the channel matrix might not reflect realistic channels well. An alternative would be to adopt one of the well-known wireless path loss models (like for example the log-distance model with a given path loss exponent) and to consider random node positions. With such a channel model, especially when the path loss exponent is three or larger, most of the channels will be either very good or very bad, whereas the range of distances leading to intermediate values will be small and only sparsely populated with relayers. In such a setting it would then also become meaningful to compare the random placement of relayers against a situation in which relayers are placed *optimally*, e.g. there could be one separate relayer for each source, placed in the middle between the source and the central controller.

References

- [1] A. Willig, "Recent and Emerging Topics in Wireless Industrial Communications: A Selection", *IEEE Trans. Industrial Informatics*, vol. 4, no. 2, pp. 102–124, May 2008.
- [2] HART Communication Foundation, *TDMA Data Link Layer Specification, HCF SPEC 075 Revision 1.1, 17 May, 2008*.
- [3] HART Communication Foundation, *HART Communication Protocol Specification, HCF SPEC 13 Revision 7.1, 05 June, 2008*.
- [4] S. N. Diggavi, N. Al-Dhahir, A. Stamoulis, and A. R. Calderbank, "Great Expectations: The Value of Spatial Diversity in Wireless Networks", *Proceedings of the IEEE*, vol. 92, no. 2, pp. 219–270, Feb. 2004.
- [5] K. J. R. Liu, A. K. Sadek, W. Su, and A. Kwasinski, *Cooperative Communications and Networking*, Cambridge University Press, Cambridge, UK, 2009.
- [6] G. Kramer, I. Maric, and R. D. Yates, "Cooperative Communications", *Foundations and Trends in Networking*, vol. 1, no. 3-4, pp. 271–425, 2006.
- [7] J. N. Laneman, D. N. C. Tse, and G. W. Wornell, "Cooperative Diversity in Wireless Networks: Efficient Protocols and Outage Behaviour", *IEEE Trans. Information Theory*, vol. 50, no. 12, pp. 3062–3080, Dec. 2004.
- [8] A. Scaglione, D. L. Goeckel, and J. N. Laneman, "Cooperative Communications in Mobile Ad Hoc Networks", *IEEE Signal Processing Magazine*, vol. 23, no. 5, pp. 18–29, Sept. 2006.
- [9] V. authors, "Special issue on Cooperative Networking – Challenges and Applications (Part I)", *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 2, pp. 241–492, Feb. 2012.
- [10] J. Cannons, L. B. Milstein, and K. Zeger, "An Algorithm for Wireless Relay Placement", *IEEE Trans. Wireless Communications*, vol. 8, no. 11, pp. 5564–5574, Nov. 2009.
- [11] D. Tse and P. Viswanath, *Fundamentals of Wireless Communications*, Cambridge University Press, Cambridge, UK, 2005.
- [12] A. Goldsmith, *Wireless Communications*, Cambridge University Press, Cambridge, UK, 2005.
- [13] H. Bai and M. Atiquzzaman, "Error Modeling Schemes for Fading Channels in Wireless Communications: A Survey", *IEEE Communications Surveys and Tutorials*, vol. 5, no. 2, pp. 2 – 9, 2003, <http://www.comsoc.org/livepubs/surveys>.
- [14] A. Willig and E. Uhlemann, "PRIOREL-COMB: A Protocol Framework Supporting Relaying and Packet Combining for Wireless Industrial Networking", in *Proc. 7th IEEE International Workshop on Factory Communication Systems (WFCS)*, May 2008, Dresden, Germany.
- [15] J. Banks, J. S. Carson, B. L. Nelson, and D. M. Nicol, *Discrete-Event System Simulation*, Prentice-Hall, Upper Saddle River, NJ, third edition, 2000.
- [16] G. Gamba, F. Tamarin, and A. Willig, "Retransmission Strategies for Cyclic Polling over Wireless Channels in the Presence of Interference", *IEEE Trans. Industrial Informatics*, vol. 6, no. 8, pp. 405–415, 2010.
- [17] A. Willig, "How to Exploit Spatial Diversity in Wireless Industrial Networks", *IFAC Annual Reviews in Control*, vol. 32, no. 1, pp. 49–57, Apr. 2008.
- [18] E. Baccarelli, M. Biagi, C. Pelizzoni, and N. Cordeschi, "Maximum-Rate Node Selection for Power-Limited Multi-antenna Relay Backbones", *IEEE Trans. Mobile Computing*, vol. 8, no. 6, pp. 807–820, June 2009.
- [19] A. S. Ibrahim, A. K. Sadek, W. Su, and K. J. R. Liu, "Cooperative Communications with Relay-Selection: When to Cooperate and Whom to Cooperate With?", *IEEE Trans. Wireless Communications*, vol. 7, no. 7, pp. 2814–2827, July 2008.
- [20] A. Nosratinia and T. E. Hunter, "Grouping and Partner Selection in Cooperative Wireless Networks", *IEEE Journal on Selected Areas in Communications*, vol. 25, no. 2, pp. 369–378, Feb. 2007.