

Probe-based Transmission Power Control for Dependable Wireless Sensor Networks

Wolf-Bastian Pöttner, and Lars Wolf

Technische Universität Braunschweig

Institute of Operating Systems and Computer Networks (IBR)

Mühlenpfordtstr. 23, 38106 Braunschweig

Email: [poettner|wolf]@ibr.cs.tu-bs.de

Abstract—Dependable Wireless Sensor Networks for industrial process automation, monitoring and control applications have high demands regarding network reliability in terms of timely delivery and packet loss. In large industrial facilities, interference between neighboring networks is a major problem that can be mitigated by reducing the radio transmission power to the necessary minimum. However, reliability of the network should not be compromised even in spite of packet loss bursts.

In this paper we aim at keeping the link within certain burstiness bounds while still minimizing the transmission power. We present evaluation results showing that the burstiness changes with the receiver signal strength. Based on this observation we present two mechanisms for dependable wireless sensor networks and study the performance in our testbed.

I. INTRODUCTION

Wireless Sensor Networks (WSNs) for industrial process monitoring and control are expected to provide high reliability in terms of timeliness and data delivery. Control loops which are realized using WSN technology require a high fraction of the data to be delivered in time for correct operation. Monitoring of industrial facilities also has high demands regarding delivery reliability. In general, networks for industrial process automation are based on Time Division Multiple Access (TDMA) technology to guarantee certain reliability and timing limits. One example of the underlying TDMA technology is GinMAC [1]; it has been developed within the GINSENG project [2], which studied WSN approaches for monitoring and control of oil refinery processes.

Dependable networks relying on TDMA schedules with fixed slot assignments have to make assumptions about the reliability of each link while crafting the respective schedules. Slots have to be provisioned to allow for retransmissions whenever a packet loss occurs. Only when the amount of lost packets is higher than what the schedule is designed for, end-to-end packet loss will occur. In dependable networks, losses have to be avoided to not jeopardize network operation. Research has shown that packet losses occur in bursts and that time-critical networks have to be provisioned for the worst-case burst that can occur [3]. The behavior of burst losses on a link is generally referred to as *burstiness* and has to be kept in mind while crafting a TDMA schedule.

In industrial facilities [2] the density of sensors is high. Since real-time networks such as GINSENG are limited in size (around 30 nodes), multiple independent networks have to be installed in parallel [4]. In the GINSENG target scenario, an oil refinery in Portugal, more than 1000 networks would

have to be installed to connect all 35000 sensors that are currently present. Neighboring networks interfere with each other, which may lead to packet loss. Even when using all channels available in IEEE 802.15.4, reusing channels is inevitable. Furthermore, not all channels are equally well suited for such networks because of background noise [5]. Since higher transmission power generally means more potential for interference [6], we aim at reducing the transmission power to the necessary minimum. Since the spatial coverage correlates with transmission power that translates into reduced potential interference. Wireless link conditions are changing over time and therefore the minimum amount of transmission power that is necessary to transmit packets from sender to receiver is also changing over time.

Stationary nodes in industrial facilities are likely connected to mains power, because the sensors and actuators that are installed consume more energy than batteries of reasonable capacity can provide. However, batteries may power some nodes that are installed in remote locations, connected to sensors that consume less energy or that are mobile. For those devices, reducing the transmission power also reduces the energy use of the node as the radio is the major contributor to energy use on sensor nodes [7] and eventually prolongs the lifetime on a set of batteries. Therefore, we pursue two goals: Minimizing interference while reducing energy consumption at the same time.

We employ a Transmission Power Control (TPC) approach that is able to control the transmission power in such a way that a certain reception signal strength is achieved. However, since TDMA schedules are designed for a certain burstiness of a link we have to ensure that the link still conforms to the burstiness that was assumed while crafting the schedule. We show that the burstiness of a link changes with changing receiver signal strength values. We furthermore present an approach to continuously measure the burstiness for multiple reception power levels to select the lowest transmission power level (corresponding to the lowest spatial coverage) that is reliable enough. We evaluate the impact on the spatial coverage as well as reliability and energy expenditure and compare the results to an approach with fixed receiver signal strength target.

The remainder of this paper is structured as follows: We present and discuss related research efforts in Section II and formulate the problem of TPC in Section III. In the following Section IV we present and discuss our approach to TPC. We present evaluation results in Section V and conclude the paper in Section VI.

II. RELATED WORK

For a long time TPC has been a research topic, especially in the field of battery-powered wireless devices with limited energy resources. In the area of WSNs and ad-hoc networks a vast number of papers have been published.

Bergamo et al. [8] have proposed an algorithm called Distributed Power Control (DPC). The sender embeds the employed transmission power into all outgoing packets and the receiver uses this information to calculate the path attenuation. By assuming symmetric links, each node can calculate the necessary transmission power towards a specific neighbor from which it has received a packet. The paper presents results that are only based on simulations in mobile ad-hoc networks.

Lin et al. [9] introduced a scheme to predict the signal power at the receiving side. The algorithm uses an initialization phase in which each node broadcasts messages at different transmission power levels. Neighbors receiving these messages give feedback to the original sender. The sender then creates a prediction function per neighbor and is then able to predict the transmission power level accordingly. If the receiver receives packets with signal strength below or above specific thresholds, it notifies the sender to adjust the prediction.

Correia et al. [10] used two different feedback mechanisms to find a suitable transmission power level per neighbor. On the one hand they monitor the number of packets that are acknowledged by the destination. If this number is above a threshold, the sender is transmitting at a power level higher than necessary and the transmission power can be reduced. The transmission power is increased if fewer packets than the lower threshold are acknowledged. On the other hand, the authors store the transmission power level in each outgoing packet. The receiver is then able to calculate the signal attenuation to that specific node. When assuming symmetric channel characteristics, each node is able to calculate the amount of transmission power that is necessary to transfer data to each neighbor.

Hackmann et al. [11] used the idea of implicit feedback using acknowledgement packets. They use a sliding window to keep track of outgoing packets. If the number of acknowledged packets is below a threshold, the transmission power is increased. If the number is above a threshold, the power is reduced and the algorithm goes into a “trial” state. If the number of unacknowledged packets increases, the transmission power is increased again and the sliding window is flushed.

Xiao et al. [12] proposed an iterative approach in which they use acknowledgement packets to get feedback about the Receiver Signal Strength Indicator (RSSI) at the receiver. If the RSSI drops below a threshold, the sender doubles the transmission power. If it is above a threshold, the sender reduces the transmission power by subtracting a constant. The values are smoothed using an Exponentially Weighted Moving Average (EWMA). The multiplicative increase and additive decrease is borrowed from the congestion control methods of Transmission Control Protocol (TCP).

In a more recent paper from Pianegiani [13], another iterative approach that judges the channel quality based on Packet-Acknowledgment Reception Rate (PARR) is described. When the PARR falls below a threshold, the power is increased iteratively, otherwise decreased. Optionally, additional mea-

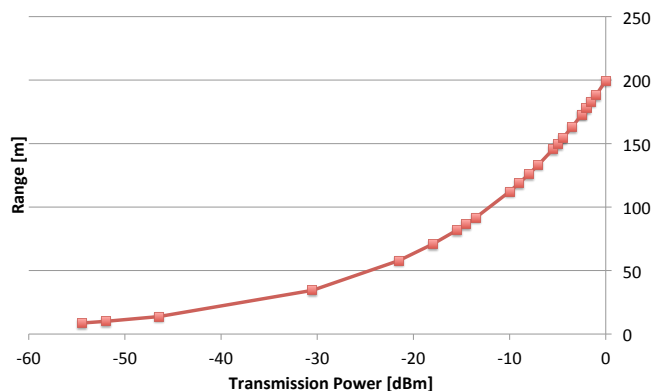


Fig. 1: Estimated range over transmission power for a receiver sensitivity of -92 dBm

asures to rate the channel quality such as RSSI or Link Quality Indicator (LQI) can be used.

The discussed publications do not respect the burstiness of links and are therefore not usable for dependable networks. Some even sacrifice reliability for additional transmission power reductions. Others assume symmetric channel characteristics or stable links that cannot be assumed in harsh environments such as oil refineries. Lastly, some publications are iterative approaches that are slow when reacting to sudden link changes. We conclude that none of the existing solutions is suitable for dependable networks.

III. PROBLEM STATEMENT

TPC mechanisms reduce the transmission power to the necessary minimum. The common goal is to save energy whereas the primary concern of this work is to reduce the interference with neighboring networks. For unsynchronized networks operating on the same channel, reducing interference can be achieved by minimizing the transmission power. The necessary TX power is the minimum transmission power that leads to successful packet reception at the receiver. However, packets that are transmitted with a power level that is too low cannot be received by the receiver and are lost. Lost packets jeopardize the network operation because delayed packets are not acceptable in dependable WSNs. If a packet gets lost and has to be retransmitted multiple times, it may not reach its destination in time. Furthermore, retransmitting packets costs additional energy.

Common radio transceiver chips like the Chipcon/TI CC2420 [14] allow selecting one out of a number of transmission power levels. The lower the transmission power, the lower the achievable range and also the energy expenditure of the chip. Figure 1 shows the estimated range based on the two-ray ground reflection model [15] with a transceiver height of 1 m and a receiver sensitivity of -92 dBm. The transmission power values have been obtained by Boano et al. in [16] and the receiver sensitivity was chosen based on extensive measurements using the CC2420 radio chip. We learn that lowering the transmission power produces a significant reduction in range and hence potential interference.

Research has shown that packet losses occur in bursts [17]. This may lead to a seemingly high statistical Packet Reception

Rate (PRR) of 99.99% whereas still 10 packets in a row get lost due to bursts. The burstiness of a link can be quantified by B_{min} and B_{max} as published by Munir et al. [3]. The authors argue that the burstiness for a given link can be expressed by the maximum amount of losses in a row (B_{max}) and the minimum number of packets that can be successfully transmitted in a row after a loss (B_{min}).

As argued in the Introduction, schedules for real-time TDMA networks have to be crafted with an underlying knowledge about the burstiness of each link to cope with loss bursts. That also means that a link that is subject to TPC has to stay within the burstiness bounds allowed by the schedule for the network to fulfill its duties. Unfortunately, the “grey zone” [18] may lead to increased packet loss because for certain reception powers it is not quite clear if a packet will be received or not. For dependable networks, links have to stay out of the “grey zone” to avoid unpredictable packet loss. Furthermore, characteristics of links change gradually over time and TPC has to adapt to such changing conditions.

We conclude that the primary goal has to be to send packets at the transmission power that achieves the lowest receiver signal strength that still leads to successful packet reception and fulfills the burstiness assumptions on which the TDMA schedule is based. Transmitting one outgoing packet at transmission power level N lets the radio transceiver consume a current of I_N at a potential V . Outgoing packets take the time t to be transmitted. Transmitting a single packet consumes $I_N \cdot t \cdot V$ of energy. A packet x is transmitted at transmission power level N_x and needs M_x transmission attempts. We can now calculate the energy e that is necessary to successfully transmit packet x with $e = t \cdot V \cdot M_x \cdot I_{N_x}$. The secondary goal of TPC is to minimize this function while at the same time ensuring that all packets arrive their destination in time.

A. Relationship between Receiver Signal Strength and Burstiness

The burstiness of a link can be measured using a mechanism called BurstProbe presented by Brown et al. [19]. Essentially, a burst of unicast packets is sent to a neighbor and the Acknowledgment (ACK) responses are collected. The probe responses allow to calculate B_{min} and B_{max} which represent the burstiness of a link.

On a given radio link, the transmitting nodes sends data with a transmission power P_t . On the path, the signal is subject to attenuation A so that the receiver signal strength P_r can be calculated as $P_r = P_t \cdot A$. According to measurements (see Section III), P_r has to be higher than -92 dBm for the CC2420 radio chip. We assume that the quality of a link (i.e. the burstiness) correlates with P_r . The intuitive explanation is that with a decreasing P_r the packet losses on a link increase and so does burstiness.

To characterize the relationship between P_r and the burstiness, we conducted extensive measurements in a testbed in an oil refinery in Portugal. In a network with 13 T-Mote Sky nodes we have measured the burstiness on all 156 links between those nodes at all transmission power levels available to the CC2420 radio transmitter. We have sent a burst of 40 probes to the respective receiver and recorded the ACK responses. We have conducted 15 measurement runs per link and power level and collected a total of 2 995 200 probe responses. We assume

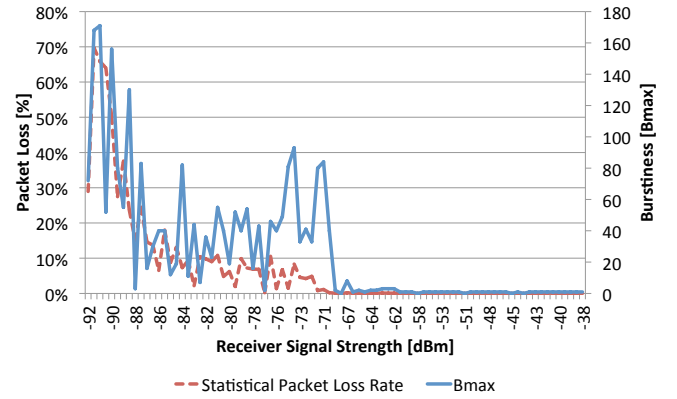


Fig. 2: Statistical packet loss rate and Burstiness for all links over receiver signal strength.

that the minimum transmission power at which at least one ACK was received corresponds to a receiver signal strength of -92 dBm. We calculate the attenuation of the link and assign a receiver signal strength to each probe pattern. We assume that the attenuation of a link remains rather constant over one measurement run and exclude links for which the standard deviation of the attenuation over multiple runs is higher than 5%.

We calculate the statistical packet loss rate by comparing the number of unacknowledged probes with the total number of probes. We plot the result for all probes on all links for each receiver signal strength in Figure 2 (individual links are shown in Figure 3). Aggregating data of many links that exhibit different characteristics causes the fluctuations. The packet loss rate increases with decreasing receiver signal strength as expected. We calculate the burstiness for the same data set and plot B_{max} over the receiver signal strength in the same figure. A higher B_{max} means that more packets have been lost in a row. We see that for a receiver signal strength from -61.0 dBm to -37.5 dBm the loss bursts are short. Most of the time, only one loss in a row occurs whereas for several receiver signal strength values no losses at all happen. In the transitional zone with a receiver signal strength from -68.5 dBm to -61.0 dBm more losses occur. With B_{max} values of up to 8, those receiver signal strength levels would require to provision a significant amount of retransmission slots. Finally, from -92.0 dBm to -68.5 dBm short loss bursts are the exception and long bursts are the rule. Those receiver signal strength levels are unusable for real-time TDMA networks.

We conclude that the burstiness increases with decreasing receiver signal strength values. This claim is supported by the fact that the statistical packet loss rate also increases with lower signal strength as expected. We further conclude that due to fluctuations a model function for the burstiness cannot be given. This means that the burstiness has to be measured for all receiver signal strength levels that could be used in order to ensure that the burstiness stays within the bounds of the TDMA schedule.

To detail our findings, we present the burstiness of three selected links in Figure 3. For link 1 we see a constantly low burstiness of $B_{max} = 0$ for receiver signal strength values from -68.0 dBm to -37.5 dBm. For lower signal

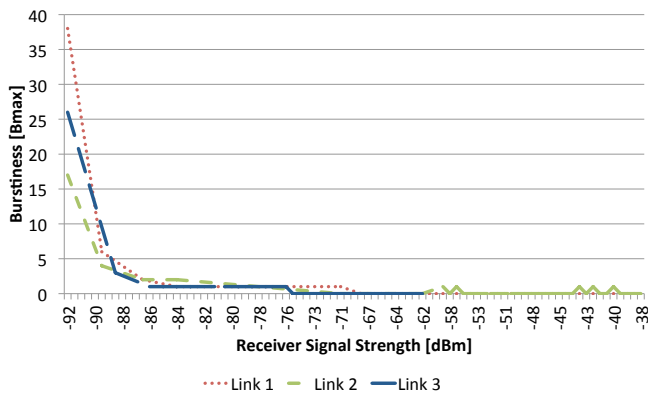


Fig. 3: Burstiness for 3 selected links over receiver signal strength.

strength values, the burstiness increases up to a maximum of 38 unacknowledged probes in a row. For link 2 we see minor fluctuations between receiver signal strength values from -37.5 dBm to -70.5 dBm with B_{max} values between 0 and 1. Finally for link 3 the burstiness is low from -61.5 dBm to -75.5 dBm with no losses and then increases up to $B_{max} = 26$ for low signal strength values. The maximum receiver signal strength for link 3 is -61.5 dBm because this link has a higher attenuation than the others.

IV. TRANSMISSION POWER CONTROL ALGORITHMS

In this section we outline our TPC approaches. We first discuss how a certain receiver signal strength target can be met using an approach to constantly measure the link attenuation. We then detail how the target can be determined by measuring the burstiness of links.

A. Attenuation-based Transmission Power Control

The concept of this approach is to control the transmission power such that a certain receiver signal strength target T is met at the receiver as published in [8]. We have extended the existing approach by employing immediate receiver feedback in the link-layer ACK packets.

We use the transmission power P_t and the receiver signal strength P_r to calculate the path attenuation A according to $A = \frac{P_t}{P_r}$ (see Section III-A). We then calculate the transmission power P_t that is necessary to produce a receiver signal strength of T at the receiver according to $P_t = \frac{T}{A}$. To feed P_r back to the sender of a frame, we embed feedback into the link-layer ACKs that allows the sender to calculate the attenuation. Using the link-layer ACKs for this purpose produces no overhead in terms of timing or energy and allows a quick response to link variations. On the sender, we calculate the EWMA of the attenuation to iron out outliers and short-term link variations. This approach allows short-term reactions to link variations as the sender receives instant feedback for each outgoing packet. Since the feedback in link-layer ACK frames only produces feedback for packets that are acknowledged, we assume P_r of unacknowledged packets to be -100 dBm. This allows reacting to packet loss by increasing the transmission power.

Compared to the published approach, our attenuation-based TPC has two advantages. Instead of assuming symmetric links,

our approach also functions on links that exhibit different characteristics in both “directions”. Furthermore, our approach uses link-layer ACK feedback and allows quick reactions to link variations without inducing any overhead. Unfortunately, this approach cannot guarantee certain burstiness for a link because it is solely based on the receiver signal strength.

B. Determining the Receiver Signal Strength Target based on Probes

Based in our findings in Section III-A we assume that the burstiness of a link changes with its receiver signal strength. Unfortunately, probing a link at various transmission powers (and corresponding receiver signal strength) values takes time and would not allow short-term reactions to link changes. Therefore, we have to determine a medium-term receiver signal strength value that is suitable using probing. This value is then used as the receiver signal strength target T as discussed in the previous section and illustrated in Figure 4.

The idea is to measure the burstiness for various receiver signal strength values and to determine the lowest value for which the burstiness is within the bounds for which the employed TDMA schedule is provisioned. This value is then used as receiver signal strength target T . However, the sender cannot influence the receiver signal strength directly. The node can only modify the transmission power while the attenuation is out of control for the sender. Therefore, not necessarily all receiver signal strength values can be probed.

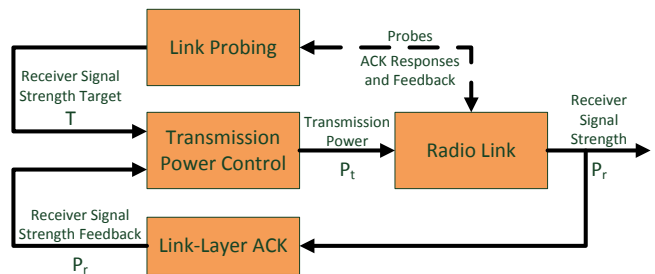


Fig. 4: Architecture for hybrid TPC approach

In order to measure the burstiness of a link, we use the concept of probe slots. Such slots are appended to the usual TDMA epoch as shown in Figure 5. In each TDMA epoch, one node exclusively uses all probe slots. We refer to this epoch as “probing epoch” of the active node. The ability to use the probe slots is distributed in a round-robin fashion among all nodes. In the probing epoch, the active node randomly picks a transmission power and uses all probe slots to send unicast packets to the upstream neighbor. The node records a pattern of acknowledged and unacknowledged packets and calculates the B_{min} and B_{max} values. P_r is fed back to the sender using link-layer ACK feedback. The sender then applies an EWMA over the P_r of all probes received by the parent. We assume that the attenuation of the link is relatively constant over the probing period (up to 15 slots of 10 ms each) and calculate B_{min} and B_{max} that are specific to this particular link at this particular P_r value and store those in a tuple R : $R = (P_r, B_{min}, B_{max})$.

The necessary amount of probe slots in each epoch depends on the burstiness a schedule can handle. For a typical schedule with a worst-case burstiness of 1/1, a reasonable number of probe slots is less or equal to 15 [19]. Those probe slots have




Fig. 5: TDMA schedule with probe slots

to be appended to the usual TDMA epoch and reduce the sleep time for the active node as well as its upstream neighbor.

All tuples R containing the burstiness for different P_r values are kept in a ring buffer. This allows invalidating old data and overwriting it with newer data to handle long-term link or hardware variations. After each probing epoch, the node iterates through all stored tuples and determines the lowest suitable P_r . Suitable means that the burstiness of the tuple is better or equal to the burstiness that is the basis for the current TDMA schedule. If multiple tuples for one P_r exist, we use the tuple representing the most lossy (and bursty) link behavior to handle worst-case situations. All tuples that are as good or better than the B_{min} and B_{max} assumptions are considered usable. We then look for the lowest P_r value among the eligible combinations and use this as the receiver signal strength target T for the attenuation-based TPC. This makes sure that the minimum specified reliability for the schedule can be met to enable safe and stable operation of our dependable WSN. It further ensures that the network stays outside of the “grey zone” because links operating in the “grey zone” expose a higher burstiness than the TDMA schedule was designed for.

Probing one link at one transmission power level takes one epoch. With an epoch length of 1 s, and $n = 13$ nodes in the network we have to probe $(n - 1) = 12$ upstream links (the sink node does not have an upstream link). With a radio that offers $m = 32$ power levels, one complete probe cycle takes $n \cdot m = 384$ s. While more probe slots offer more information about the current channel situation, they also use up more time in the epoch. For networks with tight time budget, only few slots can be allocated to probing. To produce results with higher resolution, multiple probe patterns can be concatenated. However, each individual probing period has to be long enough to capture worst-case bursts for which the TDMA schedule is provisioned.

V. EVALUATION

The goal of this evaluation is to assess the impact of the presented TPC approaches on the reliability of the network and to show the benefits in terms of interference reduction. Furthermore, effects on the energy consumption shall be quantified.

A. Evaluation Setup

We have implemented the TPC approaches discussed in the previous section for the GinMAC [1] TDMA protocol. The implementation is based on Contiki OS [20] and specifically designed for the CC2420 radio chip. We correct the errors of the RSSI readings of the CC2420 by using the approach taken by Chen et al. [21] and convert RSSI into receiver signal




Fig. 6: Evaluation testbed: 14 T-Mote Sky Nodes with GinMAC Virtual Tree Topology.

strength according to the CC2420 data sheet [14]. To provide link-layer feedback, we have used software-generated link-layer acknowledgments as published in [22]. This allows to piggyback up to 1 byte of data onto the link-layer ACKs.

We have used a total of 14 nodes with distances of less than 1 m in the basement of a university office building as displayed in Figure 6. Although unlikely, interference with active IEEE 802.11 hardware and other devices in the 2.4 GHz band cannot be excluded or quantified. While this setup appears to be simple, it is in fact a challenging situation. Nodes that close together will use one of the lowest transmission power levels. But these low levels are very prone to exterior influences and the “grey zone” is immanent. We therefore expect that mechanisms that perform well in our experiments show similar performance in the real world. Furthermore, this testbed allows conducting repeatable experiments in a rather controlled environment. To make the results more descriptive, we have run between 17 and 28 measurement runs per implementation of TPC to minimize the impact of temporal fluctuations of the environment and the radio channel and present the results in a statistically aggregated form. We plot the median as well as the minimum, maximum and the 25 and 75 % quartiles for a total of 94 experiments each lasting 3 hours and 10 minutes.

For our experiments, we have used GinMAC in a static topology configuration with a sampling frequency of 1 Hz and a maximum end-to-end delay of 1 s. The schedule contains 1 retransmission slot per transmission slot to allow links with a B_{min}/B_{max} of 1/1 to work well. Upstream traffic in GinMAC is mainly unicast while downstream traffic is broadcast. TPC only influences the transmission power that is used for unicast upstream packets towards the parent of each node. To measure the energy consumption of nodes, we have used Contiki online energy estimation feature [23] combined with current consumption values from the CC2420 data sheet [14].

In the following evaluation, we compare three implementations of TPC with baseline measurements. We refer to GinMAC without TPC as “no TPC” and use the term “TPC Attenuation” for the TPC approach with static receiver signal strength target T (see Section IV-A). Based on the results in Section III-A, we have selected a conservative receiver signal strength target $T = -60$ dBm to avoid high burstiness.

We have combined the attenuation-based TPC with the




Fig. 7: Retransmissions in the network for different implementations of TPC.

probe-based concept to determine a suitable receiver signal strength target (see Section IV-B) and refer to the combination as “TPC Hybrid”. We configure it to use a B_{min} and B_{max} target of 1/1, because the employed TDMA schedule is designed for this very value. For “TPC Hybrid” we use two different settings with either 4 or 8 probe slots at the end of each epoch. While a higher number of probe slots allows a more accurate assessment of a specific power setting, a lower number of probe slots produces less overhead in terms of time and energy consumption.

B. Impact on Reliability

To evaluate the impact of TPC on reliability, we use the number of retransmissions in the network as metric. GinMAC retransmits unicast packet if the sender of a packet does not receive a link-layer ACK in time. Therefore, the number of retransmissions counts occasions, in which the transmission power was too low to produce a successfully received data packet at the receiver. Situations in which the ACKs are lost are also possible and cannot be distinguished.

We have summed up the number of retransmissions for all nodes in the network over the whole experiment duration and present the results in Figure 7. The blue bar shows the median of all measurement runs while the black box shows the 25 % and 75 % quartiles. The whiskers show the absolute minimum and maximum. As expected, GinMAC in baseline configuration without TPC has the lowest number of retransmissions. Also, this number shows that even at the highest transmission power, still interference occurs and leads to packet loss. With a total of 159 600 regular (excluding retransmissions) transmissions in the network, only 0.08 % of the transmissions have to be repeated.

TPC Attenuation has a significantly increased retransmission count compared to “no TPC”. As expected, this increase is caused by links that expose a higher packet loss rate for lower receiver signal strength values. The number of retransmissions of TPC Hybrid with 4 probe slots is higher than the retransmissions of TPC Attenuation. This increase can be explained by the fact that TPC Attenuation uses a fixed receiver signal strength target that was chosen in a conservative way. The hybrid variant approaches the limit more aggressively and therefore also has a higher probability of packet loss. Increasing the number of probe slots from 4 to 8




Fig. 8: Total RF energy consumption of all nodes (excluding Sink node).

also increases reliability. This was expected, since using more probe slots produces a more complete picture of the channel properties. Even with the increased number of retransmissions the maximum ratio of retransmissions is quite low at 0.17 %.

C. Impact on Spatial Coverage

To evaluate the impact on spatial coverage, we use the 2-ray ground reflection model [15] to estimate the average range of nodes given the selected transmission power levels. We have used the average output power levels of all nodes to calculate the resulting range. We assume a node height of 1 m above ground and a minimum receiver sensitivity of -92 dBm. The results are stated in Table I.

A node at maximum transmission power covers an area of 199.5 m in its vicinity. Although the model may not reflect reality due to the environment in a real-world deployment, it gives a clear indication. Also, the interference range of nodes is usually assumed to be higher than the actual radio range [24]. The attenuation-based approach lowers the range to 46.3 m, which is a reduction of 76.79 %. However, the transmission power can be reduced even further as shown by the TPC Hybrid results. With a resulting radio range of 34.6 m, the spatial coverage has been reduced by 82.7 % compared to no TPC and by 25.3 % compared to the attenuation-based TPC approach. For the spatial coverage, increasing the probe slots from 4 to 8 increases the range by 2.6 % while realizing a higher reliability.

D. Impact on Energy Expenditure

To evaluate the impact on the energy expenditure, we have summed up the RF energy consumption of all nodes in the

TPC Approach	Avg. P_t	Range
No TPC	1.0000 mW	199.5 m
TPC Attenuation	0.0029 mW	46.3 m
TPC Hybrid 4	0.0009 mW	34.6 m
TPC Hybrid 8	0.0010 mW	35.5 m

TABLE I: Average Transmission Power (P_t) and estimated Node Range for TPC approaches.

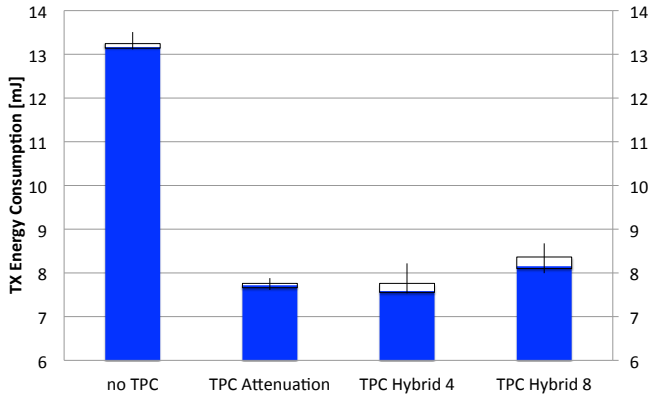


Fig. 9: Energy spent on transmitting for all nodes (excluding Sink node).

network. We have explicitly excluded the consumption of the sink node, since it has to be powered by mains power by design and is never battery-powered. Although TPC mainly influences the fraction of energy that is used to transmit data, probing approaches send more packets and therefore also increase the energy consumption for receiving data. This is why we look at the total energy expenditure instead of only the TX consumption. We expect that all variants of TPC will lower the total energy expenditure.

In Figure 8 we can see that the baseline GinMAC configuration that always transmits at the highest possible transmission power consumes a median of 42.48 mJ in the experiment duration. “TPC Attenuation” consumes a median of 37.85 mJ during the experiment and thus attributes for energy savings of 10.9% compared to “no TPC”.

“TPC Hybrid” with 4 probe slots consumes a median of 39.01 mJ for the experiment duration with an average of 39.13 mJ. The consumption is significantly higher than for attenuation-based TPC, which is expected. The probe slots for the hybrid approach consume energy for transmitting and especially for receiving and acknowledging the probes. This additional overhead is dependent on the number of probe slots and is therefore expected to increase even further for more probe slots. The results show that doubling the probe slots from 4 to 8 increases the overall energy consumption to 41.58 mJ as median or 41.67 mJ on average. Still the energy consumption with 8 probe slots is lower than the baseline measurement without TPC.

Figure 9 shows that “TPC Attenuation” and “TPC Hybrid” with 4 Probe slots all have almost identical TX energy consumption. This is interesting because the hybrid approach transmits significantly more packets when sending additional probes but is able to chose a lower transmission power so that in total, the TX energy is almost identical. The hybrid approach with 8 probe slots has a higher consumption because the 4 additional probe slots cost additional energy and cannot be compensated by even further reductions of the transmission power.

E. Discussion of the Results

The evaluation of TPC approaches for dependable networks has shown that significant potential for reduced spatial coverage exist. By using a probe-based approach to determine

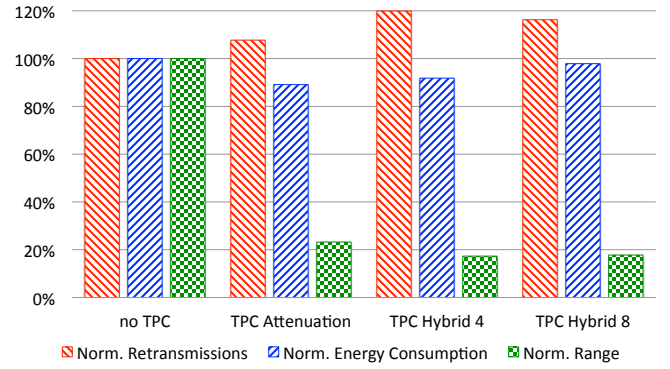


Fig. 10: Normalized comparison of the evaluation metrics.

a suitable receiver signal strength target T the covered area around a node can be reduced to a range of 36 m, which is significantly lower than the value without TPC. However, this reduction comes at increased retransmissions from 123 without TPC to 143 with TPC Hybrid. Furthermore, saving 41.37% of the energy used for transmitting data packets or 10.9% of the total energy consumption can prolong the lifetime of the network significantly.

In Figure 10, we compare the three important metrics: Retransmissions, total energy expenditure and range. We have normalized all values to GinMAC without TPC. Given the initial goal of lowering interference, the figure shows that in a lab setting “TPC Hybrid” makes the best trade-off between all three factors. “TPC Hybrid” lowers the spatial coverage to the lowest level while increasing the retransmissions only slightly. From an energy-savings point of view, “TPC Attenuation” is the best choice because it saves the maximum amount of energy. The advantage of “TPC Hybrid” over “TPC Attenuation” is that the former does not need a specified RSSI target value T and is solely based on B_{min} and B_{max} values that are the basis for TDMA schedules for dependable networks. Also, it is expected that the hybrid approaches will perform better than the other approaches in dynamic environments in which the properties of the radio channel change unexpectedly and in which readjustment of parameters is important.

The experiments with 4 or 8 probe slots illustrate that more probe slots consume more energy. However, more probe slots also achieve a more reliable network. This implies that to adequately assess the conditions and the burstiness of a radio channel, 8 probe slots are better suited than 4. It can be expected that using more probe slots will improve the reliability even further. Finally, the experiments have confirmed that TPC can help reducing the spatial coverage significantly. Depending on the application and optimization objective, any of the three evaluated TPC settings can be appropriate.

VI. CONCLUSIONS

Reducing spatial coverage in real-time WSN targeted onto industrial settings is crucial, because many networks have to operate in parallel and interference has to be avoided to keep guarantees regarding timeliness and reliability. In this paper we have established that existing TPC approaches are not well-suited for dependable networks, because the algorithms usually have a negative impact on reliability especially in the light of burst packet losses.

We have experimentally verified that the burstiness of a link increases with lower receiver signal strength values. This is an important observation because it enables us to determine receiver signal strength targets for TPC by measuring the burstiness for specific receiver signal strength values.

Based on this fact, we have presented our approach for TPC in dependable networks that measures the link attenuation using instant receiver feedback and calculates the appropriate transmission power. Furthermore, we use continuous link probing to dynamically adjust the receiver signal strength target T to cope with changing conditions. The resulting target T is based on the B_{min} and B_{max} metric that is the basis for TDMA schedules in dependable networks. Our hybrid approach is self-configuring and does not need a predetermined (and fixed) receiver signal strength target T . It is solely based on the B_{min} and B_{max} parameters used to create the TDMA schedule.

We have presented extensive experimental results regarding three TPC settings in a testbed. The results show significant potential for reduced spatial coverage. Since reduced spatial coverage correlates with interference range and thus with interference, this is an important step to increase reliability in large-scale high density dependable WSN deployments. Nevertheless we have seen that all TPC approaches increase packet loss by a certain margin. However, all TPC approaches discussed in this paper have a significant impact on spatial coverage while the negative impact on reliability is only marginal.

ACKNOWLEDGMENT

This work has been partially supported by the EC under the FP7 contract ICT-224282 (GINSENG).

REFERENCES

- [1] P. Suriyachai, J. Brown, and U. Roedig, "Time-Critical Data Delivery in Wireless Sensor Networks," in *Distributed Computing in Sensor Systems*, ser. Lecture Notes in Computer Science, R. Rajaraman, T. Moscibroda, A. Dunkels, and A. Scaglione, Eds. Springer Berlin / Heidelberg, 2010, vol. 6131, pp. 216–229.
- [2] W.-B. Pöttner, L. Wolf, J. Cecilio, P. Furtado, R. Silva, J. Silva, A. Santos, P. Gil, A. Cardoso, Z. Zinonos, J. do O, B. McCarthy, J. Brown, U. Roedig, T. O'Donovan, C. Sreenan, Z. He, T. Voigt, and A. Jugel, "WSN evaluation in industrial environments first results and lessons learned," in *Proceedings of the 4th International Workshop on Performance Control in Wireless Sensor Networks (PWSN 2011) in conjunction with IEEE DCSS 2011*, June 2011, pp. 1–8.
- [3] S. Munir, S. Lin, E. Hoque, S. M. S. Nirjon, J. A. Stankovic, and K. Whitehouse, "Addressing burstiness for reliable communication and latency bound generation in wireless sensor networks," in *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*, ser. IPSN '10. New York, NY, USA: ACM, 2010, pp. 303–314.
- [4] T. O'Donovan, J. Brown, F. Büsching, A. Cardoso, J. Cecilio, J. do O, P. Furtado, P. Gil, A. Jugel, W.-B. Pöttner, U. Roedig, J. sa Silva, R. Silva, C. Sreenan, V. Vassiliou, L. W. T. Voigt and, and Z. Zinonos, "The GINSENG System for Wireless Monitoring and Control: Design and Deployment Experiences," *ACM Transactions on Sensor Networks (TOSN)*, vol. 10, no. 3, Aug. 2014, accepted for publication.
- [5] T.-D. Tran, R. Silva, D. Nunes, and J. Sa Silva, "Characteristics of Channels of IEEE 802.15.4 Compliant Sensor Networks," *Wireless Personal Communications*, vol. 67, pp. 541–556, 2012.
- [6] J. Jeong, D. Culler, and J.-H. Oh, "Empirical analysis of transmission power control algorithms for wireless sensor networks," in *Networked Sensing Systems, 2007. INSS '07. Fourth International Conference on*, June 2007, pp. 27–34.
- [7] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*, ser. SenSys '04. New York, NY, USA: ACM, 2004, pp. 95–107.
- [8] P. Bergamo, A. Giovanardi, A. Travasoni, D. Maniezzo, G. Mazzini, and M. Zorzi, "Distributed power control for energy efficient routing in ad hoc networks," *Wirel. Netw.*, vol. 10, no. 1, pp. 29–42, Jan. 2004.
- [9] S. Lin, J. Zhang, G. Zhou, L. Gu, J. A. Stankovic, and T. He, "Atpc: adaptive transmission power control for wireless sensor networks," in *Proceedings of the 4th international conference on Embedded networked sensor systems*, ser. SenSys '06. New York, NY, USA: ACM, 2006, pp. 223–236.
- [10] L. H. A. Correia, D. F. Macedo, A. L. dos Santos, A. A. F. Loureiro, and J. M. S. Nogueira, "Transmission power control techniques for wireless sensor networks," *Comput. Netw.*, vol. 51, no. 17, pp. 4765–4779, Dec. 2007.
- [11] G. Hackmann, O. Chipara, and C. Lu, "Robust topology control for indoor wireless sensor networks," in *Proceedings of the 6th ACM conference on Embedded network sensor systems*, ser. SenSys '08. New York, NY, USA: ACM, 2008, pp. 57–70.
- [12] S. Xiao, A. Dhamdhere, V. Sivaraman, and A. Burdett, "Transmission power control in body area sensor networks for healthcare monitoring," *IEEE J.Sel. A. Commun.*, vol. 27, no. 1, pp. 37–48, Jan. 2009.
- [13] F. Pianegiani, "miniTP: A protocol for the minimization of the transmit power in wireless networks," in *Instrumentation and Measurement Technology Conference (I2MTC), 2011 IEEE*, May 2011, pp. 1–6.
- [14] Texas Instruments Incorporated, "2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver," <http://www.ti.com/lit/gpn/cc2420>.
- [15] T. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2001.
- [16] C. Boano, N. Tsiftes, T. Voigt, J. Brown, and U. Roedig, "The impact of temperature on outdoor industrial sensor network applications," *Industrial Informatics, IEEE Transactions on*, vol. 6, no. 3, pp. 451–459, Aug. 2010.
- [17] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin, "Temporal properties of low power wireless links: modeling and implications on multi-hop routing," in *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*, ser. MobiHoc '05. New York, NY, USA: ACM, 2005, pp. 414–425.
- [18] J. Zhao and R. Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *Proceedings of the 1st international conference on Embedded networked sensor systems*, ser. SenSys '03. New York, NY, USA: ACM, 2003, pp. 1–13.
- [19] J. Brown, B. McCarthy, U. Roedig, T. Voigt, and C. J. Sreenan, "Burstprobe: debugging time-critical data delivery in wireless sensor networks," in *Proceedings of the 8th European conference on Wireless sensor networks*, ser. EWSN'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 195–210.
- [20] A. Dunkels, B. Grönvall, and T. Voigt, "Contiki - a Lightweight and Flexible Operating System for Tiny Networked Sensors," in *Proceedings of the First IEEE Workshop on Embedded Networked Sensors (Emnets-I)*, Tampa, Florida, USA, Nov. 2004.
- [21] Y. Chen and A. Terzis, "On the mechanisms and effects of calibrating rssi measurements for 802.15.4 radios," in *Wireless Sensor Networks*, ser. Lecture Notes in Computer Science, J. Silva, B. Krishnamachari, and F. Boavida, Eds. Springer Berlin / Heidelberg, 2010, vol. 5970, pp. 256–271, 10.1007/978-3-642-11917.
- [22] W.-B. Pöttner, S. Schildt, D. Meyer, and L. Wolf, "Piggy-Backing Link Quality Measurements to IEEE 802.15.4 Acknowledgements," in *Mobile Adhoc and Sensor Systems (MASS), 2011 IEEE 8th International Conference on*, Oct. 2011, pp. 807–812.
- [23] A. Dunkels, F. Osterlind, N. Tsiftes, and Z. He, "Software-based online energy estimation for sensor nodes," in *Proc. of EmNets '07*. New York, NY, USA: ACM, 2007, pp. 28–32.
- [24] G. Zhou, T. He, J. A. Stankovic, and T. Abdelzaher, "RID: Radio Interference Detection in Wireless Sensor Networks," in *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2005)*, vol. 2. IEEE Computer Society, Mar. 2005, pp. 861–901.