# EchoRing: A Low-Latency, Reliable Token-Passing MAC Protocol for Wireless Industrial Networks

Christian Dombrowski
COMSYS Group
RWTH Aachen University, Germany
dombrowski@comsys.rwth-aachen.de

James Gross
School of Electrical Engineering
KTH Royal Institute of Technology, Sweden
james.gross@ee.kth.se

***Abstract*—Spurred by recent industrial trends, such as factory automation or phase synchronization in the smart grid, there is a significant interest for wireless industrial networks lately. In contrast to traditional applications, the focus is on carrying out communication at very short latencies together with high reliabilities. Meeting such extreme requirements with wireless networks is challenging. A potential candidate for such a network is a token-passing protocol, as it allows to bound latencies. However, it lacks mechanisms to cope with the dynamics of wireless channels. In this paper, we present EchoRing, a novel wireless token-passing protocol. Cooperative communication and an improved fault tolerance allow this decentralized protocol to support industrial applications over wireless networks. Based on experimental results, we demonstrate the suitability of EchoRing to support demands of industrial applications. EchoRing outperforms other schemes by several orders of magnitude in terms of reliability for latencies of and below 10 ms.**

## I. Motivation

Wireless industrial networks have gained increasing interest in the research community over the last years [1]. After initially considering wireless systems for sensing and monitoring applications in the industrial context, the focus currently shifts towards realizing critical control loops over wireless networks. Application scenarios can be found in the area of factory automation, but also in associated areas, like robotics. Using wireless connections allows control processes to be more flexible and less costly in terms of maintenance. It even enables entirely new application scenarios, e. g., the integration of moving robots and surveillance infrastructure into time-critical control processes. The associated requirements are novel as comparably small packets need to be exchanged with very short latencies *and* high reliabilities. Typical industrial low-latency applications require latencies of and below 20 ms [2]. In addition, standards like IEC 61508 [3] define Safety Integrity Levels (SILs), which are accepted quantitative metrics for the reliability of industrial plants, including their control. Regulatory bodies demand industrial processes to be compliant with a certain SIL. For instance, in case of SIL 2, reliabilities of at least 1E-6 of the entire system (including the communication) need to be guaranteed. Therefore, the underlying communication systems usually follow a redundant, partially autonomous, and decentralized organization to reach the required reliability levels.

Realizing these reliability, latency, and flexibility requirements over wireless networks is challenging. Degradations in transmission quality are very dynamic as subtle changes in the propagation field may have severe impact on the transmission. On top, industrial environments pose additional challenges due to electromagnetic interference and unfavorable materials [4]. A further challenge arises from the contradictory nature of latency and reliability. Protocols use the diversity of various domains to correct transmission errors, whereas time is the most common one. However, as latencies become shorter, the channel variability might not be high enough to allow for exploiting time diversity. As fault-tolerance is a crucial characteristic of industrial networks, other diversity sources need to be considered. This allows the networks to react quickly to link state changes to ensure reliability [5].

These challenges make it complex to design a solution [5], [6] that satisfies stringent industrial application requirements. A specifically designed set of protocols is required; in particular an appropriate Data Link Layer (DLL). Several industrial fieldbus systems [7], [8] use an approach that directly addresses the latency and flexibility requirements: Token-passing offers strict latency bounds and deterministic medium access. Yet, maintaining the stability of a token-passing network over erroneous wireless links becomes challenging [9], which leads to claims [10] that token-passing is unsuitable for wireless industrial networks. On the other hand, the decentralized and self-configuring approach of token-passing is advantageous in principle. Centralized networks suffer from severe outages if wireless channel degradations isolate critical stations.

This paper focuses on token-passing protocols and addresses the question of how to achieve stability in such a network, given a fluctuating wireless channel. Stability in the context of this paper refers to the occurrence of changes in the network topology. We propose to realize this stabilization through cooperative communication, as well as additional fault-tolerance, and refer to the resulting protocol as *EchoRing*. We argue that for token-passing protocols both approaches come at almost negligible overhead. This refutes concerns [11] that non-contention-based DLL protocols suffer from complex channel management when using cooperative communication. The second major contribution of this paper is an *experimental evaluation* of EchoRing and several comparison schemes, based on the Wireless Open-Access Research Platform (WARP) prototyping environment [12]. It demonstrates that both proposed mechanisms increase network stability and allow to achieve

remarkable packet loss rates of 1E-5 and better at latencies of 10 ms in a real-world example setting, despite raw errors rates of up to 1E-1. This shows the feasibility of applying token-passing in wireless industrial networks to realize the latency, reliability, and flexibility demands.

Industry and academia suggest various candidates for application in wireless industrial networks [6], [13]. However, only few of them address the low-latency regime, e. g., WISA [14] or RNP [15]. The majority of the candidates is based on central coordinators, which in turn limits flexibility and introduces a single point of failure. Yet, their development indicates strong interest in realizing time-critical processes over wireless networks. Notable decentralized candidates are TrustedWireless, or extensions to IEEE 802.11. A white paper on TrustedWireless [16] states the usage of a semi-decentralized approach to form wireless industrial networks in automation scenarios. Due to the lack of details of how to coordinate critical data traffic over varying network topologies and of how to react to changes in the environment, it is unknown whether this technology provides the required reliability specifically at short latencies. On the other hand, studies [17] show that also IEEE 802.11e is inappropriate for reaching low-latencies and high reliabilities in industrial environments.

This paper is structured as follows. In the beginning, Sec. II introduces token-passing and its main deficiencies in the wireless domain. Sec. III presents the design of EchoRing. The experimental measurement campaign is presented and evaluated in Sec. IV, while Sec. V concludes this work.

## II. Brief Overview of Token-Passing Protocols

Token-passing is a Medium Access Control (MAC) scheme which has been standardized for various tethered networks, e. g., IEEE 802.4 or IEEE 802.5. Today, it mainly finds application in industrial fieldbus systems, e. g., ProfiBus [7]. The main principle relies on granting medium access to a station by an exclusive transmission right. This right is forwarded among the stations by means of a token packet. The reception of the token is acknowledged. Although the physical structure is a token-bus system (inherited from the broadcast nature of the wireless medium), the stations form a logical ring. Each station has a predecessor and a successor. It is up to the protocol variant and the physical topology whether stations only have this information or an overview of the entire ring. Every station in a ring holds the token for a specified amount of time, called Token Holding Time (THT). Within the THT, a station may transmit actual payload packets including acknowledgments or exchange structural control packets. The Target Token Rotation Time (TTRT) represents the worst-case duration it takes a token to reach the station again. It is given by the product of THT and the number of stations in the ring. The TTRT is immediately the worst-case latency bound of the system for an unidirectional transmission. This assures each station a minimum data rate, which also holds for high utilization.

As the ultimate goal is to carry out a reliable payload transmission, the protocol has to deal with external influences. Transmission errors or station failures constitute the two main problem sources. If not dealt with, both problems ultimately lead to critical events called *ring instability* or *payload packet loss*, respectively. A stable ring is a precondition for carrying out time-critical payload transmissions Hence, the protocol needs to resolve a failed token forwarding process to the successor effectively to prevent ring instabilities. Automatic Repeat reQuest (ARQ) can be employed to correct the first problem, i. e., token transmission errors. Nevertheless, retransmissions may still fail as they are subject to the THT. It is up to the protocol of how to proceed. In case of ProfiBus, a station forwarding the token detects a token loss by a missing acknowledgment within a certain time. Once a specified number of retransmissions cannot resolve the situation, the station holding the token excludes its unresponsive successor and tries to forward the token to the successor of the unresponsive station. To do this, each ProfiBus station needs to maintain a valid list of all operating stations in the network. If excluded, a station subsequently only has limited real-time capabilities as it has to rejoin the ring at a later point in time. The more long-lasting the error pattern or the more stations are affected by it, the more likely is the impairment of real-time capabilities of *all* stations. Successively removing stations from the ring may lead to the worst-case scenario: An entire ring disintegration. Then, it takes considerable time to build up a ring from scratch.

The second cause of a token loss occurs if the station holding the token fails, i. e., the token disappears completely. As a result, a timer at another designated station expires. This station continues by creating a new token and starting a new token rotation cycle. Thus, stations residing in between the failed and the designated station temporarily lose real-time capabilities as they are overtaken. The previously introduced exclusion strategy that follows an unsuccessful ARQ routine removes the failed station in a second step.

Given a station is part of a stable ring, payload transmission themselves may also fail. A payload packet loss results. ARQ alleviates the effects of transmission errors. Once the end of the THT is reached, the station defers further payload transmissions to the next THT.

Overall, as error handling becomes tedious, a general requirement is the need for relatively stable links, which is the case in tethered networks. However, bringing token-passing to the wireless domain is likely to cause more problems since the wireless medium differs in three distinctive aspects from tethered networks. First, in comparison to wired links, errors occur much more often. Error patterns in wired scenarios either last long due to broken cables or station outages, or are infrequently induced by noise, whereas wireless suffers from error bursts caused by fading effects. Corresponding mitigation schemes, e. g., ARQ, are not efficient for low latencies. Second, while the probability of multiple affected links is diminishing in tethered systems, mobility and propagation effects can cause correlated failures in multiple wireless links simultaneously. This readily leads to situations in which excluding successors results into ring disintegrations since stations suffer from outdated topology information. Third, in contrast to wired systems it is much more difficult for

wireless systems to differentiate between a station failure and a wireless link degradation. Tethered systems may allow to check for basic connectivity with different means than the ongoing transmission, e. g., an idle signal of an optical fiber indicates whether a connection to a successor is still available.

Consequently, known token-passing schemes fail to cope with the wireless channel. Conservative retransmission values allow token-passing to be used in wireless scenarios [18]. Still, transmission error bursts significantly affect the ring stability [9], which ultimately diminishes the real-time characteristics of the protocol.

## III. Design of EchoRing

This section presents the main contribution of this paper, EchoRing, an extended token-passing protocol that copes with unstable wireless channels. Conventional recovery schemes lose effectiveness since errors happen more frequently on potentially more links simultaneously. These schemes exclude stations too often and too easily. It is advisable to prevent errors early on but also to identify erroneous situations quickly to prevent exclusions and potentially resulting ring disintegrations, yet let actions happen not too early to allow for error recovery. If stations have the additional capability to differentiate between transmission errors and station failures, they can choose an appropriate strategy, e. g., by keeping the successor in case of a temporal transmission error.

EchoRing achieves much more robust behavior due to the two main mechanisms of cooperation and a fault-tolerant recovery process. Sec. III-A provides details on how to lower the chances of payload packet and token losses by means of cooperation. Advanced failure recovery procedures allow for a more stable network operation if a token is lost despite using cooperation, as presented in Sec. III-B. Both mechanisms require certain protocol design decisions regarding the general error handling strategy. The implications of these decisions will be introduced in Sec. III-C.

### A. Cooperative Communication

The first robustness enhancing strategy takes advantage of cooperative communication. Cooperative ARQ allows to select another station dynamically to let it assist in the transmission process. The aim is to increase packet delivery rates and, thus, improve network stability. The task of the cooperating station, referred to as relay, is to retransmit a packet in case of a missing acknowledgment, i. e., to *echo* it. Cooperating stations have to be determined carefully to not adversely affect the success of the transmission process. Cooperative ARQ benefits mostly from spatial diversity [1], [19]. Utilizing other diversity sources is required since the wireless medium lacks time diversity on the envisioned very short latencies. An immediately retransmitted packet is very likely lost as well due to temporal correlation. Numerous work has been done on theoretical benefits [19], but also on cooperation strategies that address the questions of when and how to cooperate [11], [20], [21]. Yet, related work does not address the low-latency regime, i. e., the overhead of obtaining, distributing, and processing the required network information on short time scales may be prohibitive. Moreover, related work [11] identifies complexity issues when bringing cooperation techniques to allocation-based MAC protocols.

In EchoRing, every station employs cooperative ARQ autonomously from the others due to the decentralized approach. Hence, stations need Channel Quality Information (CQI) for as many connections in the network as possible. CQI serves as a basis for the cooperation scheme, but also means overhead in terms of acquisition. Stations exchange CQI pro-actively instead of acquiring it on-demand. This allows to achieve shorter latencies but produces considerable control traffic in the network. However, exploiting token transmissions to determine and distribute CQI causes *no overhead* in terms of additionally sent packets. This is a key feature of EchoRing. The frequently exchanged token replaces the need of other techniques to exchange explicit control packets [11]. Regarding CQI determination, EchoRing utilizes signal strengths of all received and overheard packets to obtain Signal-to-Noise Ratio (SNR) estimates describing the incoming links. This SNR estimate represents the sole CQI in order to reduce computational complexity. Yet, stations require CQI about their outgoing links and other connections in the network in order to determine the best cooperating station. Hence, the periodic token exchange is also used to distribute the required information. Every station has a local connectivity matrix containing the SNR estimates for each directional connection in the ring. All stations put CQI about their incoming links as digital data into transmitted packets, while other stations use this information to update their local connectivity matrix. As the TTRT is in general around a few milliseconds, the resulting CQI matrix can be expected to be quite accurate, yet not fully instantaneous. Although it is advantageous that stations overhear all other transmissions, it is no requirement as it does not prevent general applicability of cooperative ARQ.

The decentralized approach of maintaining a local connectivity matrix allows each station to select the best cooperating station individually. EchoRing allows to select at most one cooperating station, which still allows to achieve full diversity order [22]. In principle, different approaches exist to determine the best relay [11]. As a first step, we decide to employ an outage capacity model [23] as the basis for the relay selection scheme. The goal of using the outage capacity model is not to predict the current level of reliability, but rather to determine the best relay by assuming that the determined relay will be the optimal regardless of a specific metric. Apart from being computationally manageable, it does not require perfect CQI, i. e., fully instantaneous values, and it further allows to add CQI smoothing without loss of applicability. The outage probability model quantifies outages, i. e., error rates, based on average channel SNR and a requested capacity by extending the Shannon Capacity concept. Considering the transmission of a packet of specific size within an associated time frame leads to the requested capacity. A long-term channel capacity can be derived by using the average channel SNR. However, the actual capacity typically differs from the long-term capacity

and, hence, may not be large enough to support the requested capacity at all times. Assuming a Rayleigh-distributed block-fading channel allows to make statements about the likeliness of such an outage event. The expected total error probability of a cooperative ARQ transmission, $P^{\star}_{SD}(R)$, builds upon link outage probabilities $P_{(\cdot,\cdot)}$ belonging to connections between a given source $S$, destination $D$, and cooperating station $C$. This total error probability is given by

$$P^{\star}_{SD}(C) = P_{SD} \cdot P_{SC} + P_{SD} \cdot (1 - P_{SC}) \cdot P_{CD}.$$

For a set of available stations, the source needs to select that particular one minimizing the error probability. The set of potentially cooperating stations $\mathcal{C}(S, D)$ is defined as

$$\mathcal{C}(S, D) = \{\text{station} \in \text{local conn. matrix}\} \backslash \{S, D\}$$

Considering one relay, the optimization problem is

$$\begin{aligned} & \min_{C \in \mathcal{C}(S,D)} \; P^{*}_{SD}(C) \qquad (1) \\ \Rightarrow \; & \min_{C \in \mathcal{C}(S,D)} \; P_{SD} \cdot (P_{SC} + (1 - P_{SC}) \cdot P_{CD}) \\ \Rightarrow \; & \min_{C \in \mathcal{C}(S,D)} \; \frac{1}{\text{SNR}_{SC}} + \frac{1}{\text{SNR}_{CD}}, \end{aligned}$$

in which $\text{SNR}_{(\cdot,\cdot)}$ denotes the respective average channel SNR. Under the assumption that all systems are identical, i. e., employ the same Physical Layer (PHY) and hardware platform, the optimal choice is independent of all parameters but the average SNR. Algorithms can solve Eq. (1) with manageable efforts in $O(n)$, i. e., they are linear in the number of available cooperating stations. This allows a very fast execution appropriate for low-latency systems.

As cooperation is not always optimal, each station chooses periodically for every entry in its connectivity matrix whether relays are available and which relay is optimal, or if the station should retransmit the packet itself. The station holding the token delegates the ARQ action to a previously determined cooperating station. Each packet contains the designated relay. This is sufficient since if the packet would miss to reach the relay, it would not have a packet to send anyway. If the relay misses this packet, no station retransmits. The only additional overhead of cooperative ARQ is a by 3 B increased packet.

### *B. Evolved Failure Tolerance Mechanisms*

Besides the robustness introduced by cooperative ARQ, a station can also benefit from deliberately overhearing transmissions between other stations in its vicinity. Reliability enhancing protocol procedures allow to detect certain error events retrospectively, which enables to recover from a failure event. Yet, slight changes in the error handling strategy of token-passing protocols are required to not immediately penalize an unresponsive station (see Sec. III-C).

From a pure link layer perspective, bursty transmission errors easily lead to interpreting the link state as permanently broken, when in reality it will be stable after a few token rotation cycles. If the connection to the successor is temporarily unavailable, a station holding the token may retrieve vital information later on by taking advantage of the broadcast nature of the wireless channel, Analyzing overheard packets may allow the station to move back to normal operation instead of suffering from frequent ring instabilities. A station tries to infer that a previous inability to forward the token belongs to a certain error type, based on the overheard packets. Three *recovery procedures* act as complementing failure tolerance mechanisms. Each of them targets a specific error condition:

*1) Uni-directional transmission error:* While an ARQ scheme resolves erroneous situations if the forward direction (from Sta. 1 to successor Sta. 2) is affected, a dedicated procedure for the backward direction is beneficial. If Sta. 1 misses an acknowledgment that its token indeed reached Sta. 2 successfully, it should not take any further actions. Since Sta. 2 is typically unaware of this error, it continues operation in the current ring setup[1]. Sta. 1 can recover from this error by examining other packets. It stays in the ring if it finds packets from the same ring it already belongs to.

*2) Bi-directional transmission error:* Sta. 2 creates a new token in case of such an error and refreshes the ring.[2] This time, Sta. 1 may detect other station's packets that belong to the refreshed ring. Once Sta. 1 overhears such a transmission, it continues normal operation and notifies Sta. 2 of its continuing presence once it gets the token.

*3) Station fault:* This error is caused by Sta. 2 being shut down unexpectedly or moved out of transmission range. Hence, Sta. 2 is unable to leave the ring properly, and consequently, Sta. 1 fails to reach Sta. 2 permanently. Sta. 3 (Sta. 2's successor) notices a missing incoming token and, thus, creates a new token and refreshes the ring. As Sta. 3 marks Sta. 2 as a problematic station, Sta. 1 can deduct a station fault by comparing this problematic station to its own successor. As a result, Sta. 1 bridges the connection by making Sta. 3 its new successor and notifying Sta. 3 once it receives the token.

Although the last strategy also covers transmission error patterns involving three stations within one rotation, these recovery strategies are designed to deal with erroneous situations that involve only up to two stations. It is prohibited to use these recovery procedures after a station observed a failure on its incoming link in order to avoid a prolonging phase of errors, e. g., due to bisected rings. Once a failure in the current token rotation cycle is resolved, the presented strategies can correctly fix an error involving two other stations in the next one.

### *C. Adaptation of Error Handling Strategy*

The error handling strategy used in the majority of tethered token-passing protocols undergoes a significant change in EchoRing. As it is the station's responsibility to forward the token, this station has to take further actions instead of immediately penalizing the successor. This gives the forwarding station, but also other stations, the possibility to correct this error subsequently, since it takes the token an entire rotation to return, e. g., by applying additional error handling routines

[1] Sec. III-C offers details on collision avoidance and ring management.
[2] Token loss handling is explained in Sec. III-C in detail.

(see Sec. III-B). Only if the error cannot be fixed, real-time capabilities of the forwarding station get affected in the next token rotation cycle. In case a station fails to forward the token to its successor, it decides to leave the ring the next time it gets the token instead of excluding its successor. The following rejoin happens potentially at another logical position in the ring. The station's intention of rejoining is to establish a more stable connection between its predecessor and its successor.

Since the token is lost in the forwarding process, a token recreation timer is needed. Allowing every station to recreate a token instead of only allowing designated coordinators to do this prevents affecting real-time capabilities of other uninvolved stations. The timer is set to strictly maintain the TTRT. The next still operating successor station notices its timer expiry and continues operation. By creating a new token, the station refreshes the old ring with a new ring. This new ring has the same topology and is notified by the creating station in subsequent tokens and by setting the predecessor as a problematic station. Stations that receive or overhear subsequent tokens associate to the new ring transparently.

This new error handling strategy makes EchoRing more susceptible to certain errors if a station decides to pass the token early. Immediately forwarding the token if no payload needs to be sent or passing the token before the end of the THT is not allowed. If stations pass the token early, a conflict may arise in the next rotation cycle in case of a token loss: TTRT timer expiries of different stations converge too much to ensure a complete THT for each station. This would result in packet collisions or stations being overtaken. Hence, each station holds the token for its entirely allowed THT.

## IV. Experimental Performance Evaluation

To show the credibility of the EchoRing approach, a real-world measurement campaign is conducted. Any model-based evaluation (mathematical analysis or simulation) is subject to abstraction errors, which may have a significant impact at the low latencies and high reliabilities of interest. Hence, EchoRing is experimentally evaluated in this section using a real-world WARP testbed [12]. The goal of the performance evaluation is to investigate the achievable levels of reliability at very short latencies under varying conditions. A parameter study is conducted to identify critical parameters. EchoRing is compared to other MAC schemes to make statements about the suitability of the protocols to be applied as the basis of wireless industrial networks. In addition, the impact of EchoRing's reliability enhancing features on the performance need to be identified. Hence, less complex token-passing variants lacking EchoRing's unique features are considered.

*1) Comparison Schemes and Evaluation Metrics:* Five decentralized protocols have been evaluated which may provide wireless connectivity in industrial settings. Besides **EchoRing**, two token-passing variants were considered. On the one hand, a less complex variant is chosen that lacks the feature of cooperative ARQ, referred to as **RecoveryRing**. On the other hand, a rudimentary token-passing protocol called **BasicRing** is chosen that also lacks the recovery procedures. This variant resembles the behavior of tethered token-passing systems, which immediately exclude stations in the token forwarding path, e. g., ProfiBus. In addition, two Carrier-Sense Multiple Access (CSMA) variants were adapted that came shipped with the WARP device. They do not feature an RTS/CTS scheme and have a back-off slot length of $22\,\mu$s The first variant, later referred to as **SCSMA**, has the same number of retransmissions as EchoRing. A conservative maximum contention window is chosen that increases exponentially from 32 to 256 slots. The intention is to try maintaining the required deadline at the expense of dropped packets. In the second variant, called **CSMA**, the protocol retransmits up to 8 times with a maximum contention window of 32 to 1024. The rationale behind choosing this protocol is to allow more packet retransmissions to achieve a higher packet reception rate, however at the expense of an increased delay.

TABLE I
Measurement system parameters

| Parameter | Value |
|---|---|
| Payload size | $\mathcal{D} = 100\,\mathrm{Byte}$ |
| Deadline | $\mathcal{T} = 10\,\mathrm{ms}$ |
| Traffic periodicity | $\mathcal{P} = 50\,\mathrm{ms}$ |
| Averaged PHY PER (token, data) | 1.1E-4, 1.1E-3 |
| Token/payload modulation, coding | BPSK / QPSK, 1/2 rate |
| Bandwidth and channel | $B = 10\,\mathrm{MHz}$ in empty $5.7\,\mathrm{GHz}$ band |

The WARP platform offers a PHY similar to IEEE 802.11a. Packet headers of fixed $24\,\mathrm{B}$ size, lack of a scrambling, and a $10\,\mathrm{MHz}$ bandwidth constitute main differences.

To quantify the protocol performances and decide on their suitability for industrial scenarios, we use the metrics

- *Packet Loss Rate* (PLR), quantifying lost and late payload packets despite using ARQ,
- Payload packet *Delay Distribution* (DD), describing payload packet delays using a complementary empirical Cumulative Distribution Function (ceCDF), and
- *Ring Instability* that represents the frequency of failures in the token-forwarding process of a specific station (only applies to token-passing schemes).

*2) Parameterization of Measurement Campaigns:* Various measurement campaigns were planned to investigate the performance of all schemes. The parameterization of the default scenario is given in Table I, which is in line with typical industrial parameters [2]. The intention was to see how PLRs and DDs behave under typical conditions. To investigate the system behavior at short latencies, only one retransmission of payload and token packets, respectively, is allowed per THT. Campaigns with varying ring size, Packet Error Rates (PERs), and payload traffic intensity were conducted subsequently to evaluate system behavior under even more challenging conditions. In particular, a higher PER is generated by lowering the transmission power to simulate worse channel and environmental conditions.

*3) Experimental Methodology:* The measurement topology consists of up to five WARP devices. The environment in

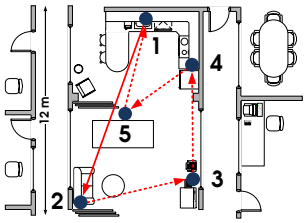


Fig. 1. Measurement environment

which the measurements are carried out is an office environment shown in Fig. 1, spanning about 12 m. Data packet exchanges are visualized with arrows.

Several measurement campaigns are conducted, each of which lasts for 60 min. Measured metrics are collected and aggregated on-board. It is assured that the aggregation does not interfere with the functionality of the MAC protocol.

Payloads of fixed size $\mathcal{D}$ are generated periodically with interval $\mathcal{P}$ and enqueued in a priority queue at each station. Generation intervals are randomly aligned. Payload transmissions happen between Sta. $x$ and Sta. $x+1$. If the duration between the events of enqueuing at the source and notifying the successful reception to higher layers at the destination exceeds the deadline $\mathcal{T}$, the packet is considered to be late. To judge on the lateness, Sta. 1 and Sta. 2 are time-synchronized, i. e., one oscillator drives both boards' logic circuits. The average PERs are obtained from the actual campaigns by considering in each THT only the first token from the predecessor and data packet from the specified source, respectively. Hence, note that PERs are not only conservative estimates (e. g., , if ring collapsed due to bad channel conditions, immediately recurring errors cannot be counted since station is out of ring), but also lack information about time correlation (larger than 80%).

Confidence intervals are given for a 95% confidence level and at least 40 replications using a student-t distribution for interval determination. All figures use a logarithmic scale for the y-axis. Despite the log scale, most y axes contain a "0" to show the case of having observed no packet errors.

### A. Protocol Evaluation for Various Stations

In a first step, all protocols are evaluated at all stations for the default parameterization. Fig. 2a presents the payload PLRs. No bar is shown if no packets were lost. The entry of Sta. 2 marked with a star indicates the payload PLR if late packets are also regarded as lost. Fig. 2a reveals that the BasicRing protocol achieves the worst performance (a PLR of 1E-3), while EchoRing achieves the best reliability with PLRs of 1E-6 or better. In fact, Sta. 2 even had no error events recorded at all. The RecoveryRing protocol achieves a performance in between, with PLRs of 1E-4 or better. This also holds for the SCSMA variant that features the same number of retransmission as the token-passing variants. Comparing RecoveryRing and SCSMA shows that adding the proposed recovery procedures compensates the structural overhead of maintaining a stable ring. The CSMA variant outperforms RecoveryRing and reaches the packet delivery level of EchoRing,

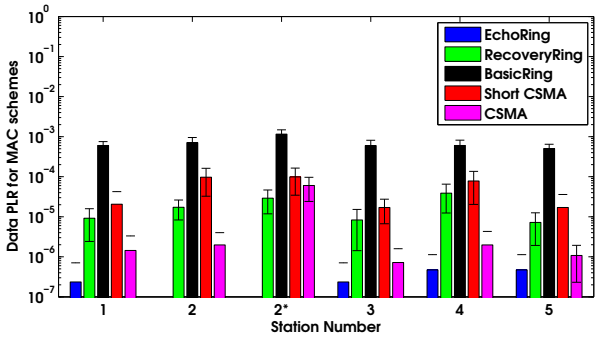


(a) Data PLRs for all stations (star marks consideration of lateness)

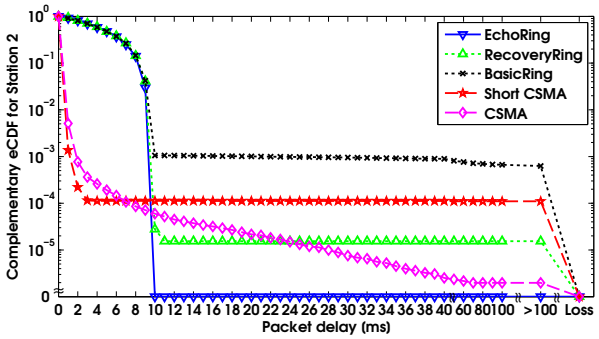


(b) Packet Delay Distribution at Sta. 2 (Note the scaling breaks)

Fig. 2. Evaluation results for a 5-station ring

however at the expense of late packets (compare both magenta-colored bars for Sta. 2 entries). Comparably long backoff periods of CSMA that result from up to eight retransmissions allow to benefit from temporal diversity to a certain degree, but cause latency violations. Note that the PLRs were achieved despite raw PERs of 1E-3 and a time correlation of $\geq 80\%$. In the remainder of this work, the focus is on Sta. 2, while treating late packets as lost.

Corresponding packet Delay Distributions (DDs) are given in Fig. 2b. The complementary empirical Cumulative Distribution Function (ceCDF) of packet delays is plotted on a log-scale, since very rare events are of particular interest. Furthermore, we added the "0" level to allow for a more intuitive interpretation. A packet loss is interpreted as delay = $\infty$. The figure shows that BasicRing has the worst performance due to many rejoin phases. Its heavy tail accounts for a considerable amount of late packets, while for SCSMA the tail is mainly dominated by packet losses (due to a rather low system load). In case of CSMA, the probability of observing a longer delay is higher, but the probability of entirely losing a packet is almost two orders of magnitude better than SCSMA. However, for EchoRing we observed no late packets or packet losses, and hence, its ceCDF quickly reaches zero. Note that for all token-passing variants, the THT is fixed, regardless of payload packets to be transmitted or not. This leads to a relatively large mass in the range of $0 - 10$ ms.

### B. Protocol Evaluation for Varying Parameters

First, the impact of the ring size on the PLRs is examined. For a ring size of four stations, Sta. 4 is switched off,

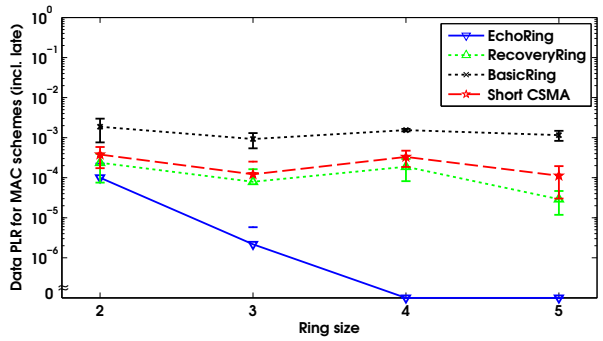


Fig. 3. Sta. 2's PLRs for various ring sizes (Note y-axis' scaling break)

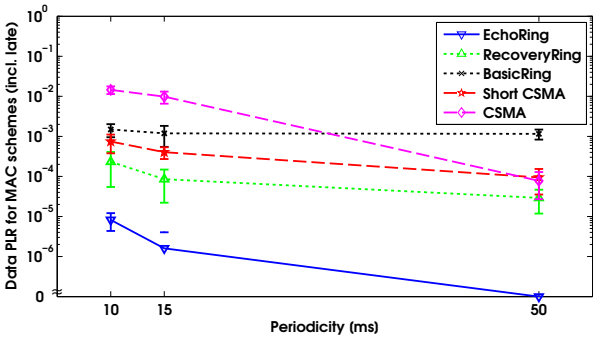


Fig. 4. Sta. 2's PLRs for various traffic loads (Note y-axis' scaling break)

while Sta. 4–5 and Sta. 3–5 for ring sizes of three and two, respectively (see Fig. 1). Fig. 3 provides the obtained PLRs. BasicRing, RecoveryRing and SCSMA are almost stable for all ring sizes (partially overlapping confidence intervals). BasicRing is again worse by one order of magnitude due to ring instabilities. SCSMA and RecoveryRing are mainly limited by raw payload PERs which can be concluded from the fact that the achievable PLRs are approximately the same, especially in case of only two active stations. EchoRing exhibits considerably better PLRs by using cooperative ARQ for an increasing ring size. The more stations are part of the ring, the more diversity potential can be realized, and thus, the more stable EchoRing becomes.

Next, we analyze the performance of all protocols under higher system load. Payload packets are not generated only every $50\,\mathrm{ms}$, but also at $10$ and $15\,\mathrm{ms}$ intervals. The PLRs of all token-passing variants stays almost constant over the evaluated periodicity range, except for the case of $\mathcal{T} \approx \mathcal{P}$. While payload transmissions may not be affected by ring instabilities in case of long $\mathcal{P}$, these instabilities affect one or more transmissions if $\mathcal{P}$ approaches $\mathcal{T}$ (recall that stations need to rejoin the ring). More notable is the loss in PLR performance for the CSMA variants, caused by the contention based medium access. As a result, the favorable delay profile of SCSMA in Fig. 2b gets heavily shifted towards higher latencies (not shown). However, all EchoRing variants do not lose their latency characteristics if the traffic intensity is not increased beyond the assigned transmission capacity.

Finally, the impact of less reliable links is evaluated. This was done by decreasing transmission power by up to $30\,\mathrm{dB}$.

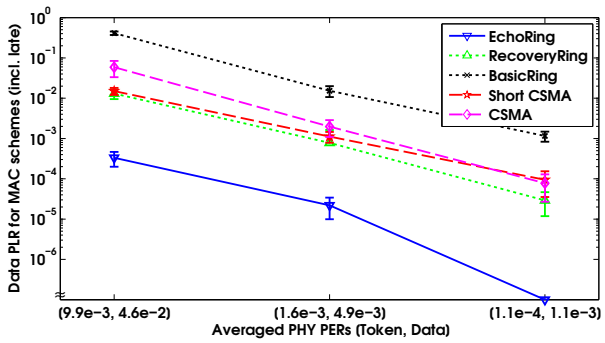


Fig. 5. Sta. 2's PLRs for various transmit powers (Note y-axis' scaling break)

While PHY payload PERs are in the range of 1E-3 for maximum power, by lowering the power average PERs increase drastically to 5E-2, while few links deteriorate even up to 2E-2. The resulting data PLRs for Sta. 2 are given in Fig. 5. As expected, raw PERs degrade as the transmission power is reduced. All token-passing variants suffer from an increased number of ring instabilities due to more errors in the token forwarding process on the one hand and payload packet losses on the other hand. This effect can be weakened significantly in case of EchoRing, as it allows to choose relays dynamically to benefit from cooperative diversity. Although both CSMA variants have a similar PLRs in case of full transmit power, their performance soon degrades due to an increased number of collisions caused by retransmissions. While SCSMA manages to achieve the reliability of RecoveryRing, CSMA fails since its back-off algorithm causes considerable delays.

### C. Evaluation of Robustness Enhancements

In the end, details are provided of how cooperative ARQ and fault-tolerant design avoid ring instabilities and improve payload PLRs. The three token-passing variants are compared for various ring sizes and the default parameterization (Table I). Table II contains detailed information on PLRs of payload packets and tokens, the absolute number of ring instabilities and the ratio of how many rotations are affected, and the additional traffic caused by cooperative ARQ.

By examining Table II it is obvious that BasicRing suffers from the lacking robustness in terms of direct packet losses and considerable ring instabilities (one in $31{,}000$ rotations, i. e., one in $5.3\,\mathrm{min}$ of continuous transmission using $10\,\mathrm{ms}$ token rotation cycles). A token loss directly translates into a ring instability. Introducing the improved recovery procedures changes this significantly. Despite slightly larger token PLRs, the number of ring instabilities is almost negligible. Obviously, this does not hold for two-station rings since no other traffic can be overheard. Payload PLRs get reduced by up to two orders of magnitude, even without applying cooperative ARQ. Adding cooperation increases the payload PLRs by another one to two orders of magnitude, up to the point of no observed losses. Utilizing cooperative ARQ contributes towards more reliable behavior by improving the payload PLR directly, and by reducing the likeliness of ring instabilities additionally (less than one instability in $5.5$ days). Note that while larger rings

TABLE II
STA. 2'S STATISTICS OF ECHORING'S ENHANCEMENTS

| Ring size | 2 | 3 | 4 | 5 |
|---|---|---|---|---|
| **BasicRing** | | | | |
| Rx Data PLR | 1.88E-4 | 9.23E-4 | 1.55E-3 | 1.16E-3 |
| Rx Token PLR | 1.45E-5 | 1.43E-5 | 2.61E-5 | 3.17E-5 |
| No. of ring instabilities | 680 | 561 | 886 | 1189 |
| Ring instability ratio | 1.45E-5 | 1.43E-5 | 2.61E-5 | 3.17E-5 |
| **RecoveryRing** | | | | |
| Rx Data PLR | 2.36E-4 | 7.94E-5 | 1.88E-4 | 2.92E-5 |
| Rx Token PLR | 3.44E-5 | 1.88E-5 | 2.31E-4 | 7.62E-5 |
| No. of ring instabilities | 145 | 0 | 7 | 11 |
| Ring instability ratio | 1.84E-6 | $<$ 1.6E-8 | 1.81E-7 | 2.83E-7 |
| **EchoRing** | | | | |
| Rx Data PLR | 1.05E-4 | 2.17E-6 | $<$4.79E-7 | $<$2.35E-7 |
| Rx Token PLR | 1.01E-5 | 3.80E-5 | 3.48E-6 | 2.90E-5 |
| No. of ring instabilities | 78 | 4 | 0 | 0 |
| Ring instability ratio | 1.02E-6 | 9.73E-8 | $<$ 3.58E-8 | $<$ 2.14E-8 |
| Add. Data Relay Traffic | – | 99.92% | 155.03% | 193.75% |
| Add. Token Relay Traffic | – | 99.97% | 141.42% | 210.22% |

get more instable in general due to an increasing probability of having bad connections in the token forwarding path, cooperative ARQ not only mitigates but even reverses this behavior by taking advantage of cooperative diversity. However, this comes at the cost of additional traffic at certain stations due to relaying, e. g., Sta. 2 forwards 210% of *additional* data traffic on behalf of other stations in the ring.

## V. CONCLUSIONS

This work proposes the EchoRing protocol as a decentralized solution for hard real-time constrained communication over wireless links. We provide details of the reliability enhancing techniques, namely cooperation strategies that exploit the broadcast nature of the wireless channel. Both strategies complement each other by (1) decreasing the payload packet loss rates, (2) improving the token forwarding process, and (3) offering synergies in case that error patterns involve more than two stations. Various parameters and topologies are evaluated using a real-world testbed. Packet error rates and packet delays indicate that applying a known token-passing scheme directly without wireless-specific modifications yields a mediocre performance. However, we demonstrate that the proposed cooperative protocol extensions allow token-passing protocols to be a reliable basis of wireless industrial networks. This disproves contrary claims regarding the general suitability of token-passing [10] and complexity issues of cooperation techniques in allocation-based protocols [11]. Certain parameterizations led to no observed packet losses or missed deadlines, despite significant raw error rates on the links and long-lasting measurement campaigns. EchoRing is a promising candidate that can pave the road towards ultra-reliable low-latency wireless industrial networks, yet more investigations are required.

## REFERENCES


[1] A. Willig, "Recent and Emerging Topics in Wireless Industrial Communications: A Selection," *IEEE Trans. on Industrial Informatics*, vol. 4, no. 2, pp. 102–124, May 2008.

[2] A. Frotzscher, U. Wetzker, M. Bauer, M. Rentschler, M. Beyer, S. Elspass, and H. Klessig, "Requirements and current solutions of wireless communication in industrial automation," in *IEEE Int'l Conf. on Communications (ICC), Workshops*, Sydney, Australia, 2014, pp. 67–72.

[3] IEC, *Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems*, IEC Std. 61 508 ed.2, 04/2010.

[4] A. Willig, M. Kubisch, C. Hoene, and A. Wolisz, "Measurements of a Wireless Link in an Industrial Environment Using an IEEE 802.11-Compliant Physical Layer," *IEEE Trans. on Industrial Electronics*, vol. 49, no. 6, pp. 1265–1282, 2002.

[5] J. Akerberg, M. Gidlund, and M. Bjorkman, "Future Research Challenges in Wireless Sensor and Actuator Networks Targeting Industrial Automation," in *IEEE INDIN*, 2011, pp. 410–415.

[6] P. Suriyachai, U. Roedig, and A. Scott, "A Survey of MAC Protocols for Mission-Critical Applications in Wireless Sensor Networks," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 2, pp. 240–264, 2012.

[7] "ProfiBus." [Online]. Available: http://www.profibus.com

[8] S. Vitturi, "Some features of two fieldbuses of the IEC 61158 standard," *Computer Standards & Interfaces*, vol. 22, no. 3, pp. 203–215, 2000.

[9] A. Willig and A. Wolisz, "Ring Stability of the PROFIBUS Token-Passing Protocol over Error-Prone Links," *IEEE Trans. on Industrial Electronics*, vol. 48, no. 5, pp. 1025–1033, Oct 2001.

[10] H.-J. Korber, H. Wattar, and G. Scholl, "Modular Wireless Real-Time Sensor/Actuator Network for Factory Automation Applications," *IEEE Trans. on Industrial Informatics*, vol. 3, no. 2, pp. 111–119, May 2007.

[11] P. Ju, W. Song, and D. Zhou, "Survey on Cooperative Medium Access Control Protocols," *IET Communic.*, vol. 7, no. 9, pp. 893–902, 2013.

[12] "WARP Project." [Online]. Available: http://www.warpproject.org

[13] S. Petersen and S. Carlsen, "WirelessHART Versus ISA100.11a: The Format War Hits the Factory Floor," *IEEE Industrial Electronics Magazine*, vol. 5, no. 4, pp. 23–34, 2011.

[14] G. Scheible, D. Dzung, J. Endresen, and J.-E. Frey, "Unplugged but Connected - Design and Implementation of a Truly Wireless Real-Time Sensor/Actuator Interface," *IEEE Industrial Electronics Magazine*, vol. 1, no. 2, pp. 25–34, 2007.

[15] J. Silvo, L. M. Eriksson, M. Bjorkbom, and S. Nethi, "Ultra-Reliable and Real-Time Communication in Local Wireless Applications," in *Conf. of the IEEE Industrial Electronics Society (IECON)*, 2013, pp. 5611–5616.

[16] F. Hakemeyer, "White Paper: Trusted Wireless 2.0 - Wireless Technology for Industrial Applications," *Phoenix Contact, Division Industrial Electronics*, 2013, ION02-13.000.PR4.2013.

[17] D. Brevi, L. Pilosu, F. Fileppo, and R. Scopigno, "Viability and Guidelines for the Effective Integration of Consumer WiFi in Industrial Plants," in *IEEE Int'l Congress on Ultra Modern Telecommunications and Control Systems (ICUMT)*, Oct 2010, pp. 232–239.

[18] M. Ergen, D. Lee, R. Sengupta, and P. Varaiya, "WTRP - Wireless Token Ring Protocol," *IEEE Trans. on Vehicular Technology*, vol. 53, no. 6, pp. 1863–1881, 2004.

[19] S. Diggavi, N. Al-Dhahir, A. Stamoulis, and A. Calderbank, "Great Expectations: The Value of Spatial Diversity in Wireless Networks," *Proceedings of the IEEE*, vol. 92, no. 2, pp. 219–270, 2004.

[20] A. Ulusoy, O. Gurbuz, and A. Onat, "Wireless Model-Based Predictive Networked Control System Over Cooperative Wireless Network," *IEEE Trans. on Industrial Informatics*, vol. 7, no. 1, pp. 41–51, 2011.

[21] N. Marchenko, T. Andre, G. Brandner, W. Masood, and C. Bettstetter, "An Experimental Study of Selective Cooperative Relaying in Industrial Wireless Sensor Networks," *IEEE Trans. on Industrial Informatics*, vol. 10, no. 3, pp. 1806–1816, Aug 2014.

[22] A. Ibrahim, A. Sadek, W. Su, and K. Liu, "Cooperative Communications with Relay-Selection: When to Cooperate and Whom to Cooperate With?" *IEEE Trans. on Wireless Commun.*, vol. 7, pp. 2814–2827, 2008.

[23] E. Biglieri, J. Proakis, and S. Shamai, "Fading Channels: Information-Theoretic and Communications Aspects," *IEEE Trans. on Information Theory*, vol. 44, no. 6, pp. 2619–2692, 1998.