A Highly Reliable Ultralow-Latency Wireless Solution for Industrial Control Loops: Design and Evaluation

Georg Kail[1], Hannes Muhr[1], Janos Gila[1], Martin Schiefer[1], Reinhard Hladik[1], Markus Hofer[2], Stefan Zelenbaba[2], Thomas Zemen[2]
[1]*Siemens Technology / RF Technologies AT & Electronic Design AT, Siemens AG Österreich*, Vienna, Austria
[2]*Security & Communication Technologies, Austrian Institute of Technology*, Vienna, Austria
Email: georg.kail@siemens.com

***Abstract*—We present the results of implementing and testing a wireless solution intended for direct one-to-one replacement of communication cables in fast industrial control loops. With an ultralow cycle time of 8000 packets/s on each unidirectional link, we measured packet error rates in the range of 1 error per 1 billion packets in the presence of industrial production-hall channel conditions, achieving an unprecedented combination of high reliability and low latency in fading channels. The transceivers were built using only easily accessible and affordable components. Frequency diversity and spatial diversity are exploited through coded orthogonal frequency-division multiplexing (OFDM) modulation and 2x2 multiple-input-multiple-output (MIMO) transmission. The replacement of cables by the devised wireless solution opens up possibilities from simplified reconfigurability and reduced maintenance cost to the introduction of entirely new production processes that are currently not possible.**

***Index Terms*—URLLC, low-latency, industrial control loop, industrial WLAN, 5G**

I. Introduction

Upcoming industry 4.0 production processes [1] are characterized by a surging demand of fast reconfigurability. Enhancing the flexibility of production environments is a rapidly growing market. In many cases, machine-to-machine communication links that require a cable pose a decisive obstacle to fast and autonomous reconfiguration, by contrast to wireless communication links. Consequently, the latter are preferred wherever suitable solutions are available, which strongly depends on the specific requirements a given application defines for its communication link. In some cases, requirements on reliability, latency, throughput, or privacy can presently only be met by wired links, while effects like random channel distortion or interference and the need for handling the ensuing transmission errors have precluded wireless solutions so far. Mastering new sets of requirements that could until now not be met by wireless links therefore opens up unprecedented possibilities from simplified reconfigurability and reduced maintenance cost to the introduction of entirely new production processes that are currently not possible.

This work was supported by the Austrian Research Promotion Agency (FFG) and the Austrian Ministry for Climate Action, Environment, Energy, Mobility, Innovation, and Technology under project 858675 (UNWIRE) of the funding program Production of the Future.

Fast control loops in production systems are a prominent example of a process which until now is confined to cabled links due to its stringent requirements on both reliability and latency. For the communication between sensors, actuators, and processing units within a control cycle, ethernet-based cabled solutions that are currently available on the market offer negligible error rates and control cycle times (bi-directional data exchange) of around 100 µs: In the Siemens product PROFINET, the DRIVE-CLiQ [2] technology allows to realize control loops with a typical cycle time of only 125 µs for motion control applications with up to three sensors and three actuators. Time-triggered Ethernet by TTTech [3] offers cycle times in the range of 100 µs. The POWERLINK [4] standard allows cycle times around 200 µs. Further solutions such as EtherCAT, Sercos III, or EtherNet/IP [5] lie in the same range. Cycle times around 125 µs and near-deterministic reliability are required in applications such as state-of-the-art chipping machines and laser-cutting machines for controlling the positioning motors of cutting heads.

For wireless links, matching such benchmarks is extremely challenging, especially due to the random fading characteristics caused by multipath scattering. Measurements of radio channels in industrial production environments have shown that such channels can exhibit distinct propagation conditions, typically with dense multipath scattering due to the abundance of metallic scatterers [6], [7] and with potentially strong variations over time [8] or between different factory environments [9]. Signal processing that prevents or corrects transmission errors caused by fading may require much larger processing times than the targeted cycle time would allow. Moreover, in cases where certain multiple-access protocols like listen-before-talk or duty cycle limitations are implemented, the minimum achievable cycle time is increased even further.

As a consequence, wireless data links for industrial environments that are currently available are not designed for benchmarks in a similar range as mentioned above for wired links. Some existing solutions are based on the IEEE 802.11 standard [10] with protocols 802.11e or 802.11g. The Siemens product iWLAN (Industrial WLAN) [11], for example, can be configured such that it guarantees cycle times in the range of 16 ms. Industrial wireless connectivity is achieved by tunneling

the aforementioned PROFINET through an iWLAN link, but communication is then limited to the performance achievable by iWLAN. The EchoRing [12] system by R3COMS claims a guaranteed latency of 10 ms, while the achievable cycle time scales with the number of nodes. Various industrial wireless links are based on the IEEE 802.15.4 standard [13], e.g., Wireless HART, ISA100.11a, ZigBee, MiWi, or RF4CE [14]. Here, cycle times are above 20 ms. The standard's extension to ultra-wideband solutions (used, e.g., by Decawave [15]), while raising the achievable data rate, further increases the cycle time due to multiple-access schemes. The WISA system [16] by ABB, based on the IEEE 802.15.1 standard [17], features a cycle time of at least 4 ms.

Contributions. This paper describes the results of implementing and testing a wireless solution intended for direct one-to-one replacement of communication cables in fast industrial control loops. With a prescribed cycle time of 125 µs, i.e., 8000 packets/s on each unidirectional link, we set ourselves the challenging goal of achieving a packet error rate (PER) of $1.45 \cdot 10^{-9}$ (i.e., 1 packet error per day), within a range of about 15 m. Measurements presented in Section III show that this unprecedented combination of high reliability and low latency in the presence of production-hall channel conditions is indeed achieved. The transceivers were built using only easily accessible and affordable components.

II. System Design

Overview. The main design parameters of the devised transceiver are as follows. Ultralow latency is prescribed in the form of a cycle time of 125 µs, as mentioned above. Using a multi-node protocol that allows for 5 uplink packets and 5 downlink packets within one cycle, a total throughput of 22.6 Mbit/s can be achieved, which is an additional advantage over the IEEE 802.15.4 and 802.15.1 systems discussed above. An alternative configuration with only 1 uplink packet and 1 downlink packet can reduce the cycle time down to 25 µs. The low-latency transceiver is designed with a transmit bandwidth of 16.25 MHz, while the center frequency can be chosen between 100 MHz and 6 GHz. A communication range of 15 m and node mobility with up to 6 m/s are envisioned.

In order to mitigate the effects of the random fading channel, the devised transceiver takes advantage of diversity in two domains. Frequency diversity is exploited through orthogonal frequency-division multiplexing (OFDM) [18] with coding and interleaving in the frequency domain. Spatial diversity is utilized both on the transmitter (TX) and the receiver (RX) side through multiple-input-multiple-output (MIMO) transmission with 2 transmit antennas and 2 receive antennas. The devised transceiver implements a modified IEEE 802.11n protocol [10], with deterministic channel occupation and more robust channel estimation based on more pilots in the OFDM symbol.

The analog transceiver tasks as well as conversion between the analog and digital domains are carried out by a 2×2 MIMO transceiver integrated circuit (IC). The signal processing tasks are carried out by a programmable logic (PL) as hardware (HW) accelerator. The PL exchanges the TX and RX payloads with an ARM9 processor, the "communication CPU", which handles the communication protocol. The communication CPU interacts with the higher-level application via another ARM9 processor, the "application-side CPU". The PL and the two CPUs are combined in a system-on-a-chip (SoC).

MIMO transceiver IC. The MIMO transceiver IC provides two direct conversion radio transmitters with a fully integrated synthesizer and two direct conversion radio receivers, also with a fully integrated synthesizer. The two digital-to-analog converters (DACs) and the two continuous-time Σ/Δ analog-to-digital converters (ADCs) operate with a resolution of 12 bits at a rate of 260 MHz. In the digital domain, the IC filters and interpolates/decimates the signals with configurable finite-impulse-response (FIR) filters, yielding a sample rate of 16.25 MHz at the PL interface. Further blocks on the RX path include a low-noise amplifier (LNA), a lowpass filter (LPF), and quadrature and DC offset calibration. Automatic gain control (AGC) is implemented with 70 dB range and 1 dB resolution. The MIMO transceiver IC uses a 40 MHz crystal for clocking. It can be programmed over a serial peripheral interface (SPI). Furthermore, there are some dedicated pins for real-time monitoring and control, which are directly used by the communication CPU.

Signal processing on PL. Signal processing requires high computational effort and is time-critical due to the stringent latency requirements. It is therefore executed by the HW accelerator. Fig. 1 shows the signal processing blocks implemented in the TX (top part) and the RX (bottom part). The left end of the figure represents the payload interface towards the communication CPU. At the right end, the antennas symbolize the MIMO in-phase/quadrature (I/Q) interface towards the transceiver IC and, subsequently, the radio link. In the TX, the PL augments the payload with a cyclic redundancy check (CRC). Through encoding and interleaving, redundancy is introduced and spread across the transmit frequency band, in order to exploit frequency diversity and mitigate the effects of frequency selective fading. After quadrature amplitude (QAM) modulation using 64-QAM and the insertion of pilot symbols for channel estimation, an Alamouti encoder [19] ensures that the transmit diversity of the two TX antennas is utilized. Finally, time-frequency transformation is performed by the OFDM modulator. In the RX, synchronization is based on a downlink beacon, and frequency offsets between TX and RX are estimated based on the received OFDM symbols. OFDM demodulation then extracts the received pilots and data symbols. The Alamouti structure is first used for pilot-based channel estimation and then, after interpolation of the channel estimates, for appropriate combination and weighting of the received distorted data symbols. Through maximum ratio combining [19], the receive diversity of the two RX antennas is exploited. Soft QAM demodulation yields log-likelihood ratios, which are then deinterleaved and decoded. The validity of the decoded data is checked through the CRC.

CPUs and interfaces. Software related to the communication protocol runs on the communication CPU as "bare metal" application (without OS). It also controls certain parameters for the PL-based signal processing and monitors the performance of the communication link. The second CPU is reserved for higher-level applications. Fig. 2 shows a schematic diagram

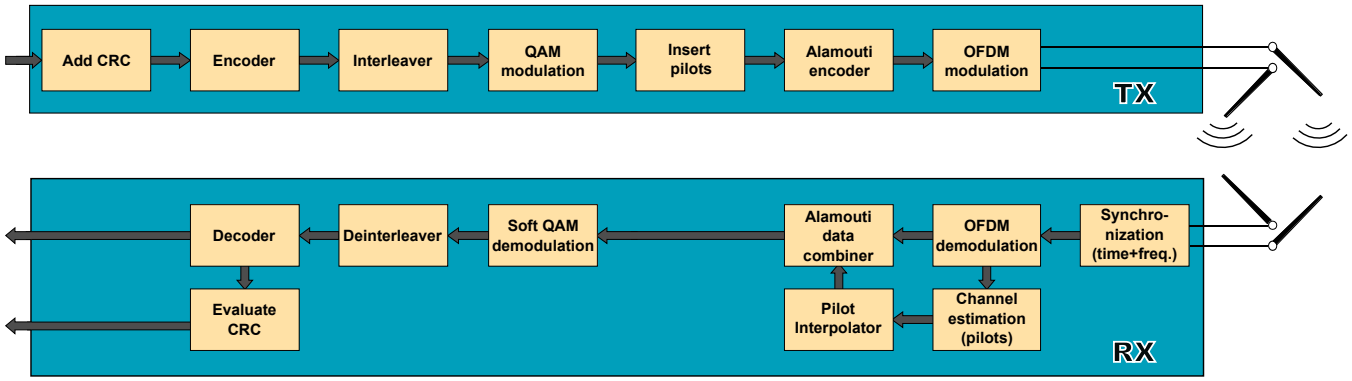

Fig. 1. Signal processing blocks between the interface towards the CPUs (left end) and the interface towards the transceiver IC and antennas (right end).

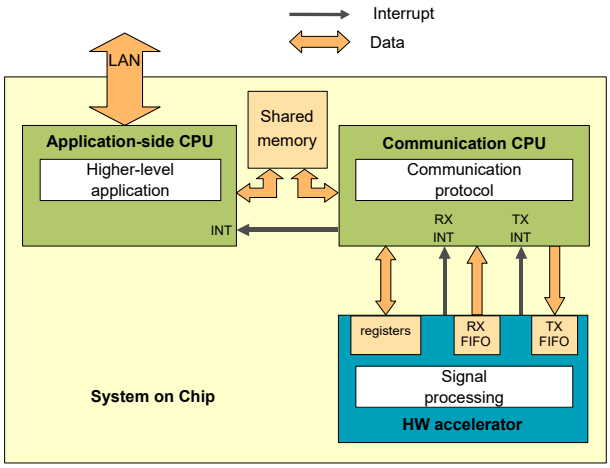

Fig. 2. Software interfaces on the SoC.

of the software interfaces on the SoC. All communication interfaces are interrupt driven to keep the latency as small as possible. The PL provides two first-in-first-out (FIFO) buffers for TX and RX frames, as well as a number of status and control registers. During HW tests the CPUs and PL are programmed either via a USB connection from a PC or via an SD card that is read during start-up. For the purpose of evaluating the radio link performance (see section III), a USB-to-UART bridge is used for outputting a log stream.

III. Measured Performance

The performance of the devised transceiver was assessed in three different ways: by numerical simulations, real over-the-air transmissions, and fading channel emulation using a HW channel emulator. The targeted PER of $1.45 \cdot 10^{-9}$ was achieved in all three approaches. Here, we present the results from channel emulation, which we find the most meaningful of the three, because this approach comprises all HW-related imperfections (by contrast to simulation) and at the same time allows us to choose challenging and representative channel properties in a controlled manner (as opposed to over-the-air transmission). Similar arguments are given in [20]. We used the newly developed AIT real-time MIMO channel emulator, which is based on the subspace emulation concept described in [21] extended for correlated 2×2 MIMO channels. During development of the devised transceivers, HW-in-the-loop tests with the channel emulator were an essential tool for reaching the challenging performance goal of the transceivers. In an iterative process, the transceivers were repeatedly improved based on new emulator measurements.

In order to assess the transceiver's performance, we recorded the PERs in emulated MIMO fading settings with different channel parameters, as described in the following paragraphs. The two TX antenna connectors of one transceiver and the two RX connectors of another transceiver are connected to the emulator. The TX continuously transmits packets with random data at a rate of 8000 packets/s. The transmission corresponds to a unidirectional link using only one of the 10 available time slots and thus achieving a raw data rate of 2.26 Mbit/s. The emulator randomly generates time-variant channel impulse responses with a specified average attenuation. The internal calculations in the emulator process the signals at a sample rate of 40 MHz. The RX CPU is configured to continuously generate a receive statistics log and stream it over its serial output. The log contains cumulative statistics for each interval of 1 s duration, most importantly the number of packets with correct CRC. With a known constant rate of packet transmissions, the PER is directly obtained from this.

For the measurements presented below, we assume interference-free transmission, i.e., a dedicated frequency band. In practice, this can be achieved through on-site management of frequency bands, especially in industrial-scientific-medical (ISM) bands. In the near future, implementation of 5G may additionally open up local frequency band licensing or local campus networks (e.g., [22] for Germany). The emulated random fading channels were generated following Clarke's model [23] with 40 paths in each channel tap. Channel taps are spaced 60 ns apart in the delay domain in accordance with the bandwidth of the transmit signal (16.25 MHz) and follow an exponentially decaying power delay profile. In order to focus on worst-case performance, only non-line-of-sight (NLOS) scenarios are considered here. Within the range supported by the transceiver IC (100 MHz–6 GHz), the transmission frequency is set to 3.4 GHz. The Doppler bandwidth is set to 68 Hz, which corresponds to a relative speed of 6 m/s. Correlation among the 2×2 MIMO subchannels is obtained by specifying a 4×4 correlation matrix. The degree of correlation in a given correlation matrix $\mathbf{R}$ can be characterized in terms

of the diversity order [24]

$$\Psi = \left(\frac{\mathrm{tr}\,\mathbf{R}}{\|\mathbf{R}\|_{\mathrm{F}}} \right)^2, \qquad (1)$$

where tr denotes the trace operator and $\|\cdot\|_{\mathrm{F}}$ denotes the Frobenius norm. Note that $1 \le \Psi \le 4$ for a 4×4 correlation matrix. In our experiments, the respective correlation matrix with integer-value Ψ is chosen such that its trace is 4 and it has Ψ nonzero eigenvalues, all of which are equal.

Measuring rates in the targeted range around 10^{-9} is a rare endeavour because of the enormous number of packet transmissions required before a certain precision can be reached. In the measurements presented here, the available resources allowed for the transmission of up to $1.1{\cdot}10^9$ packets at a few selected measurement points. With 0–2 packet errors in such an ensemble, a precise rate cannot be deducted, but it is evident that an ultralow error rate around or not far from the targeted range was achieved. The ultralow-PER results are summarized in Table I. In order to link them with the respective PER curves in higher PER domains, the ultralow-PER results are also added as text boxes in Fig. 3–5, and the measured PER curves are extrapolated (with dashed lines) into the ultralow PER domain to illustrate the qualitative message.

Table I also contains the emulated mean path loss at each of the given measurement points, i.e., the difference between the transmit power of 10 dBm and the average of the measured received signal power, which fluctuates strongly due to fading. Like the mean path loss, the mean signal-to-noise ratio (SNR) is also determined by setting the attenuation parameter of the emulator. With increasing path loss, the SNR decreases accordingly, as can be seen in the table. Table I shows that PERs in the targeted range can be achieved with a mean path loss of up to 66–70 dB (and a measured mean SNR of 30–34 dB), depending on the diversity order. A channel sounding campaign that investigated the MIMO channel between two of our transceivers in an exemplary industrial production hall setting (see [25] for details) found mean path loss values of at most 43.6 dB at a distance of up to 9.59 m, indicating that the targeted distance of 15 m should correspond to a path loss well below 65 dB and, consequently, a PER below 10^{-9}.

In Fig. 3, the effect of different degrees of correlation among the 2×2 MIMO subchannels is compared. The rms delay spread is fixed at 75 ns. The curves show that the transceivers effectively exploit MIMO diversity to lower the error rate curve, with a fairly good proportional match of the log-scale curves with the diversity order Ψ. In the absence of MIMO diversity (i.e., $\Psi = 1$), which is equivalent to a single-input-single-output (SISO) scenario, the targeted ultralow PER of $1.45 \cdot 10^{-9}$ is not achieved, since there is an error floor at about 10^{-6}. When Ψ is increased to 2, the error floor is clearly lowered so much that rates in the region of 10^{-9} are attainable. It follows that only fairly moderate MIMO diversity is required for letting the transceivers achieve the targeted PER. An analysis of the channel sounding campaign from [25] with two of our transceivers in a production hall revealed Ψ values of typically 2–2.75, with the 1st percentile at 1.53. Note, however, that these results are based on mixed line-of-sight (LOS) and NLOS scenarios. In NLOS settings as considered here, Ψ can be expected to be larger.

TABLE I
ULTRALOW-ERROR-RATE MEASUREMENTS

Parameters			Mean SNR[3]	Packet errors in $1.1{\cdot}10^9$ packets	see Fig.
TX mode	MIMO correlation[1]	Path loss[2]			
Standard	Order 2	66 dB	34 dB	2	3
Standard	Order 3	68 dB	32 dB	0	3,4
Standard	Order 4	70 dB	30 dB	0	3
Low data rate[4]	Order 3	78 dB	22 dB	0	5

[1]Diversity order is defined in (1).
[2]Mean path loss of the emulated channel, i.e., the difference between the transmit power of 10 dBm and the average of the measured received signal power (which fluctuates due to fading).
[3]The measured mean SNR scales inversely with the mean path loss.
[4]In low-data-rate mode, the data rate is reduced for the benefit of more robust coding and modulation (QPSK instead of 64-QAM).

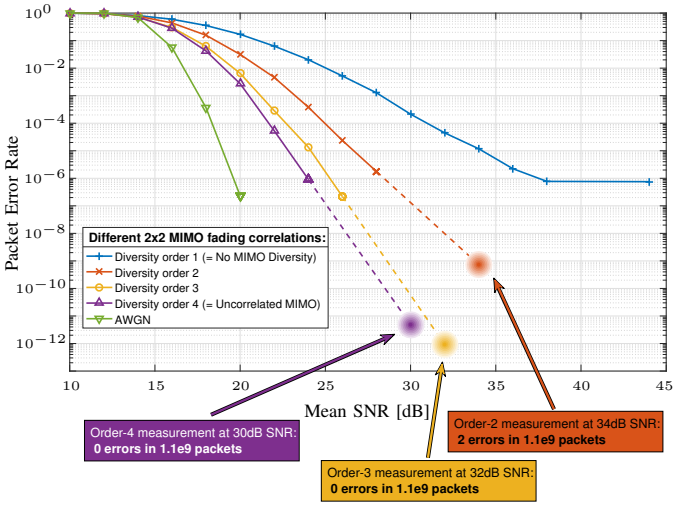

Fig. 3. Measured performance of the transceivers in Clarke's fading with rms delay spread 75 ns and Doppler spread 68 Hz. The ultralow-error-rate measurements presented in the colored boxes are associated with the corresponding curves through the dashed lines (which are extrapolations). Diversity order is defined in (1).

Fig. 4 shows how the transceivers' performance depends on the rms delay spread of the channel. The MIMO subchannels correlation is fixed at diversity order 3. It can be seen that the transceivers effectively exploit the delay diversity to lower the error rate curve. The channel sounding campaign of [25] found smaller rms delay spreads, typically in the range of 15–30 ns. However, as mentioned above, these measurements contain mixed LOS and NLOS scenarios. In NLOS settings as considered here, the delay spread can be expected to be larger. The transmit bandwidth of 16.25 MHz used by the transceivers correponds to a temporal resolution of about 60 ns, which limits the possibilities of exploiting very small delay spreads.

In Fig. 5, the performance of the transceivers in low-data-rate mode is studied. In this mode, the raw data rate is reduced to 432 kbit/s for the benefit of more robust coding and modulation (QPSK instead of 64-QAM). Again, only one of the 10 available time slots is used. As in Fig. 3, different

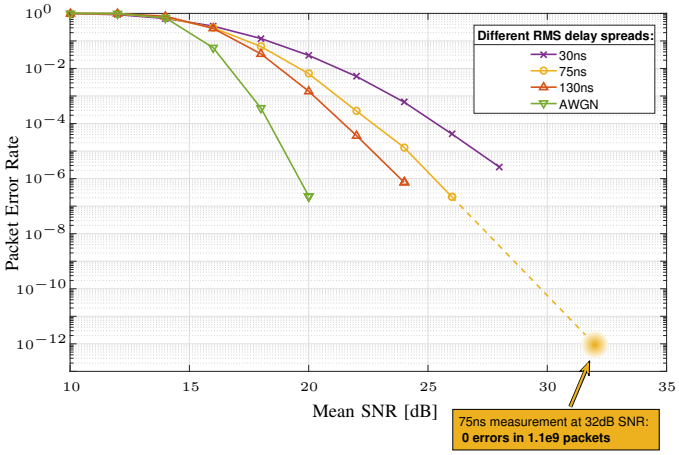

Fig. 4. Measured performance of the transceivers in Clarke's fading with Doppler spread 68 Hz. Correlation among the 2×2 MIMO channels is of diversity order 3 (cf. (1)). The ultralow-error-rate measurement presented in the colored box is associated with the corresponding curve through the dashed line (which is an extrapolation).

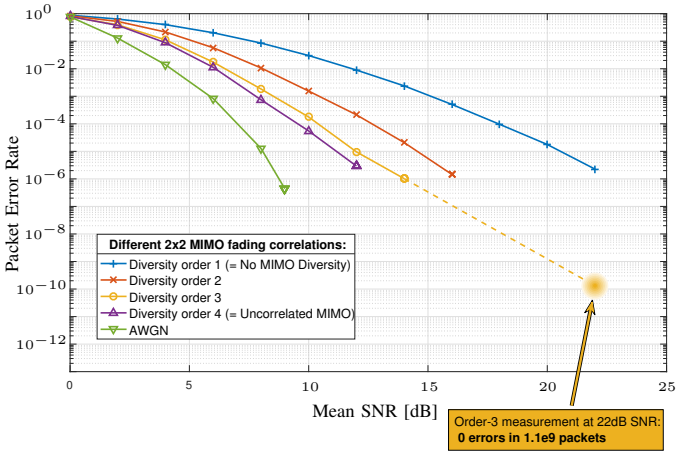

Fig. 5. Measured performance of the transceivers in low-data-rate mode in Clarke's fading with rms delay spread 75 ns and Doppler spread 68 Hz. The ultralow-error-rate measurement presented in the colored box is associated with the corresponding curve through the dashed line (which is an extrapolation). Diversity order is defined in (1).

MIMO diversity orders are compared, and the rms delay spread is fixed at 75 ns. The curves show that, compared to Fig. 3, the stronger coding and modulation lead to an additional SNR reserve of about 11–12 dB.

IV. Conclusion

We presented the system design and the measured performance of a wireless solution intended for direct one-to-one replacement of communication cables in fast industrial control loops. As required for this goal, the devised system masters the highly challenging combination of ultralow latency (cycle times around 125 µs) and ultralow error rates (PERs around $1.45 \cdot 10^{-9}$) in fading channels that resemble industrial production-hall settings. This was shown by HW-in-the-loop tests using a HW channel emulator, with settings including NLOS fading, a mean path loss of up to 66–70 dB, and node mobility with up to 6 m/s. Through measured PER curves we also showed that the devised system, which uses coded OFDM modulation and 2×2 MIMO transmission, successfully exploits frequency diversity and spatial diversity.

References

[1] K. Schwab, *The Fourth Industrial Revolution*. Crown, 2017.

[2] [Online]. Available: www.profibus.com/technology/profinet/, https://new.siemens.com/global/en/products/drives/electric-motors/motion-control-motors/motion-control-motors-accessories.html#Encoderinterface

[3] "Deterministic ethernet for high-availability, safety-, and mission-critical systems," White paper, TTTech.

[4] [Online]. Available: www.ethernet-powerlink.org

[5] [Online]. Available: www.ethercat.org, www.sercos.com, www.rockwellautomation.com/en-us/capabilities/industrial-automation-control/integrated-architecture.html

[6] J. Karedal, S. Wyne, P. Almers, F. Tufvesson, and A. F. Molisch, "UWB channel measurements in an industrial environment," in *Proc. IEEE Global Telecomm. Conf. (GLOBECOM)*, vol. 6, 2004, pp. 3511–3516.

[7] M. Cheffena, "Industrial wireless sensor networks: channel modeling and performance evaluation," *EURASIP J. Wireless Commun. Networking*, 2012, article number: 297 (2012).

[8] P. Agrawal, A. Ahlén, T. Olofsson, and M. Gidlund, "Characterization of long term channel variations in industrial wireless sensor networks," in *Proc. IEEE Int. Conf. Commun. (ICC)*, 2014, pp. 1–6.

[9] P. Stenumgaard, J. Chilo, J. Ferrer-Coll, and P. Angskog, "Challenges and conditions for wireless machine-to-machine communications in industrial environments," *IEEE Commun. Mag.*, vol. 51, no. 6, pp. 187–192, 2013.

[10] *Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, IEEE Std. 802.11, Rev. 2012.

[11] [Online]. Available: support.industry.siemens.com/cs/ww/en/view/22681042

[12] [Online]. Available: www.r3.group/echoring-technology

[13] *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low Rate Wireless Personal Area Networks (WPANs)*, IEEE Std. 802.15.4, Rev. 2006.

[14] [Online]. Available: www.fieldcommgroup.org/technologies/hart, isa100wci.org, zigbeealliance.org, www.microchip.com/en-us/products/wireless-connectivity/sub-ghz/miwi-protocol, zigbeealliance.org /solution/rf4ce/

[15] [Online]. Available: www.decawave.com

[16] R. Steigmann and J. Endresen, "Introduction to WISA," White paper, ABB Corporate Research, Jul. 2006, v2.0.

[17] *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (WPANs)*, IEEE Std. 802.15.1, Rev. 2005.

[18] J. A. C. Bingham, "Multicarrier modulation for data transmission: an idea whose time has come," *IEEE Commun. Mag.*, vol. 28, no. 5, pp. 5–14, 1990.

[19] S. M. Alamouti, "A simple transmit diversity technique for wireless communications," *IEEE J. Sel. Areas Commun.*, vol. 16, no. 8, pp. 1451–1458, 1998.

[20] J. Matai, P. Meng, L. Wu, B. Weals, and R. Kastner, "Designing a hardware in the loop wireless digital channel emulator for software defined radio," in *Proc. Int. Conf. Field-Programmable Technol. (FPT)*, 2012, pp. 206–214.

[21] M. Hofer, Z. Xu, D. Vlastaras, B. Schrenk, D. Loeschenbrand, F. Tufvesson, and T. Zemen, "Real-time geometry-based wireless channel emulation," *IEEE Trans. Veh. Technol.*, vol. 68, no. 2, pp. 1631–1645, 2019.

[22] [Online]. Available: https://www.bundesnetzagentur.de/DE/Sachgebiete/Telekommunikation/Unternehmen_Institutionen/Frequenzen/OeffentlicheNetze/LokaleNetze/lokalenetze-node.html

[23] R. H. Clarke, "A statistical theory of mobile-radio reception," *Bell Systems Technical Journal*, vol. 47, no. 6, pp. 957–1000, 1968.

[24] M. T. Ivrlac and J. A. Nossek, "Quantifying diversity and correlation in Rayleigh fading MIMO communication systems," in *Proc. 3rd IEEE Int. Symp. Signal Process. Information Technol.*, 2003, pp. 158–161.

[25] S. Zelenbaba, M. Hofer, D. Löschenbrand, G. Kail, M. Schiefer, and T. Zemen, "Spatial properties of industrial wireless ultra-reliable low-latency communication MIMO links," in *Proc. 53rd Asilomar Conf. Signals, Systems, Computers*, 2019, pp. 1054–1058.