Bringing Transmit Antenna Diversity to LPWANs: An Experimental Testbed Implementation

Felix Wunsch, Holger Jäkel, Friedrich Jondral
{FELIX.WUNSCH, HOLGER.JAEKEL, FRIEDRICH.JONDRAL}@KIT.EDU
Communications Engineering Lab, Karlsruhe Institute of Technology, Karlsruhe (Germany)

Abstract

In this paper, we present a Software-Defined Radio (SDR) testbed implementation of an experimental waveform for Low Power Wide Area Networks (LPWANs) in GNU Radio. The code is made available on GitHub. Designed specifically for the uplink in LPWANs, the waveform is based on Differential Spatial Modulation (DSM) and Direct Sequence Spread Spectrum Code Division Multiple Access (DS-CDMA). We provide a detailed overview of the testbed implementation with an emphasis on the synchronization process as well as results from simulations and over-the-air measurements in a lab environment that validate our design.

1. Introduction

Low Power Wide Area Networks (LPWANs) are an emerging paradigm for sensor networks in, e.g., Smart City contexts. In contrast to conventional cellular networks, LPWANs focus on ultra-low power operation to enable sensors to operate with a single battery for more than 10 years. Data rate requirements for popular use cases, such as monitoring and control of street lights, are usually very modest, in the range of a few bytes per hour. Mostly deployed as star-topology networks, the cell radius of an LPWAN often exceed 10 km, even in urban scenarios. This requires the devices to achieve very high processing gain to overcome distance-related path loss and provide deep indoor coverage, which is usually achieved by employing either narrowband or spread spectrum techniques (Raza et al., 2017).

While most technologies, such as LoRa, SigFox and NB-IoT, target sub-GHz frequencies due to their favorable propagation characteristics, there is also interest in utilizing the 2.4 GHz band. Compared to the unlicensed bands in the 900 MHz range, the 2.4 GHz band is globally available and does not suffer from duty cycle restrictions, which is especially important if acknowledged transmissions are desired. Even though it is well-known that this band is often congested, there is still a considerable amount of bandwidth available with acceptable interference levels (Wunsch et al., 2018b). Another advantage of higher frequencies and therefore shorter wavelengths is the more compact device size, even allowing the placement of multiple antennas at the sensor despite strict size constraints and thus opening up the possibility to exploit diversity to combat fading. Ingenu has invented Random Phase Multiple Access (RPMA) (rpm) for the use at 2.4 GHz and also actively contributed to the standardization of the Low Energy Critical Infrastructure Monitoring (LECIM) Direct Sequence Spread Spectrum (DSSS) physical layer in the IEEE 802.15.4 standard, which has recently been implemented in GNU Radio (Wunsch et al., 2017).

Multiple antennas at the sensor not only facilitate receive diversity schemes in the downlink as already employed by commercially available RPMA devices, but also transmit diversity in the uplink. Major design goals for sensors, however, are cost and power consumption, and therefore multiple transmit chains are undesirable. In (Wunsch et al., 2018a), the authors have presented an experimental waveform based on a Full Diversity Differential Spatial Modulation (FD-DSM) and DSSS, showing its superiority over existing approaches. For the sake of brevity, we will hence call it *Spatial Modulation for Long Range Sensor Networks* (*SpaRSe*). FD-DSM is a recent multi-antenna technique that enables a transmit diversity order of two with two transmit antennas and a single transmit chain.

Our contribution in this paper is that we present a real-time capable implementation of *SpaRSe* in GNU Radio that is evaluated in simulations and over-the-air measurements. The related code is made available on GitHub (Wunsch & Maier, 2018). Furthermore and to the best of the authors' knowledge, this is also the first freely available implementation of a Spatial Modulation system. The remainder of the paper is structured as follows: The details of *SpaRSe* are given in Sec. 2 followed by the description of the implementation and its design in Sec. 3. Sec. 4 provides results from simulations and over-the-air real-time transmissions. Conclusions and an outlook are given in Sec. 5.

Proceedings of the 8^{th} GNU Radio Conference, Copyright 2018 by the author(s).

2. *SpaRSe* system description

2.1. Overview

The design of *SpaRSe* is based on the IEEE 802.15.4 LECIM DSSS physical layer. It uses the same frame size, preamble structure, rate $\frac{1}{2}$ convolutional Forward Error Correction (FEC), interleaving algorithm Π, and Gold codes for spreading the preamble and the payload data, respectively. The chip modulation, however, is changed from differential BPSK (or alternatively OQPSK) to FD-DSM, which will be discussed in greater detail in Sec. 2.2. From a high-level perspective, FD-DSM can be regarded as a differential Space-Time Block Code (STBC) that encodes information over two transmit antennas and two time slots. The insertion of short time gaps between the spreading sequences is also an important alteration of the original scheme and reasoned in Sec. 2.3.

Fig. 1 shows the block diagram for a *SpaRSe* transmitter and depicts the different steps that convert the payload bits into the final transmit signal.

2.2. Full-Diversity Differential Spatial Modulation

The basic idea of Spatial Modulation (SM) is to encode information into the selection of the transmit antenna over which a conventional PSK or QAM symbol is transmitted (Mesleh et al., 2008). Only one of n_t antennas is activated at a given time, therefore transmitting an additional $\log_2(n_\mathrm{t})$ bits while (theoretically) still requiring only a single transmit chain and also enabling simple decoding using the Maximum Likelihood (ML) principle (Jeganathan et al., 2008).

In LPWANs, estimating the channel impulse reponse represents an undesirable and costly overhead, therefore differential modulation schemes are often preferred. Similarly, Differential Spatial Modulation (DSM) schemes dispense with those estimates. Additional antennas increase the degrees of freedom in the system, and while some approaches use them to increase the spectral efficiency, they can also be exploited to create transmit diversity. This is also the idea behind FD-DSM (Zhang et al., 2015), a scheme that achieves the maximum diversity order of two with two transmit antennas.

Modulation for FD-DSM can be described as follows: The input bit group $\boldsymbol{b}$ of length $k = \log_2(2L)$ selects indices $q \in \{0, 1\}$ and $\ell \in \{0, \dots, L-1\}$. The tuple (q, ℓ) selects the matrix $\boldsymbol{S}(t) \in \mathbb{C}^{2\times 2}$ which is then multiplied with the previous transmit matrix $\boldsymbol{X}(t-1) \in \mathbb{C}^{2\times 2}$ to form the new transmit matrix

$$\boldsymbol{X}(t) = \boldsymbol{X}(t-1)\boldsymbol{S}(t) = \boldsymbol{X}(t-1)\boldsymbol{A}_q\boldsymbol{V}_\ell \tag{1}$$

with

$$\boldsymbol{A}_q \in \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 0 & \exp(\mathrm{j}\phi) \\ \exp(\mathrm{j}\phi) & 0 \end{bmatrix} \right\} \tag{2}$$

and

$$\boldsymbol{V}_\ell \in \left\{ \begin{bmatrix} \exp\left(\mathrm{j}\frac{2\pi u_1 \ell}{L}\right) & 0 \\ 0 & \exp\left(\mathrm{j}\frac{2\pi u_2 \ell}{L}\right) \end{bmatrix} \right\} \tag{3}$$

where u_1 can be set to 1 and u_2 and ϕ depend on the choice of L. E.g., for $L = 2$ and therefore a spectral efficiency of $\eta = 1\,\mathrm{bit/s/Hz}$, $u_2 = 1$ and $\phi = \pi/2$ while for $L = 8$ it is $\eta = 2\,\mathrm{bit/s/Hz}$, $u_2 = 3$, and $\phi = \pi/4$, respectively. Please note that there is still only one transmit antenna active in every time slot.

Demodulation in channels with flat fading and Additive White Gaussian Noise (AWGN) is usually performed using the ML criterion, which yields

$$(\hat{q}, \hat{\ell}) = \underset{(q,\ell)}{\arg\min} \, \|\boldsymbol{Y}(t) - \boldsymbol{Y}(t-1)\boldsymbol{A}_q\boldsymbol{V}_\ell\|^2. \tag{4}$$

If soft bits, e.g., for convolutional channel decoding, are desired, these can be generated by evaluating approximate Log-Likelihood Ratios (LLRs) as

$$\mathrm{LLR}_i = \frac{\mathcal{M}_{i,1} - \mathcal{M}_{i,0}}{2\sigma_\mathrm{n}^2} \tag{5}$$

where $\mathcal{M}_{i,1}$ and $\mathcal{M}_{i,0}$ represent the smallest Euclidean distance between the received $\boldsymbol{Y}(t)$ and all possible $\boldsymbol{Y}(t-1)\boldsymbol{A}_q\boldsymbol{V}_\ell$ where the i-th bit of the bit group assigned to (q, ℓ) is 1 or 0, respectively (Martin, 2015). Dividing the distance difference by $2\sigma_\mathrm{n}^2$ accounts for the differential nature of the scheme and the reliability of the decision in presence of noise.

The authors show in (Wunsch et al., 2018a) that the introduced waveform outperforms the IEEE 802.15.4 LECIM DSSS PHY, which makes use of conventional differential BPSK or OQPSK, by up to $8\,\mathrm{dB}$ when compared at a Packet Error Rate (PER) of 1% and a payload size of 16 bytes. These results are achieved for flat and independent Rayleigh block-fading channels, but it is also stated that the scheme does not even collapse for fully correlated channels and still offers several dB of gain for highly correlated channels.

2.3. DSSS With Time Gaps

A challenging problem for practical SM systems is spectrally efficient pulse shaping. Practical pulses such as Root Raised Cosine (RRC) span multiple symbol durations and therefore switching antennas that are fed by a single transmit chain would lead to undesirable out-of-band radiation in conventional high-rate systems. This can be alleviated by employing more transmit chains as proposed in

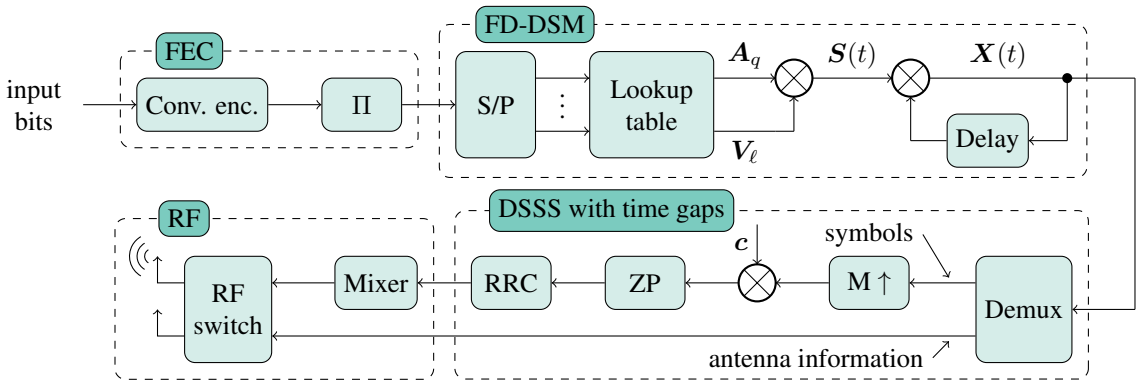

Figure 1. SpaRSe transmitter block diagram (Wunsch et al., 2018a). Zeropadding (ZP) facilitates a transmit diversity order of two with single transmit chain operation and spectrally efficient pulse shaping at the cost of only an additional RF switch. The Gold code $\boldsymbol{c}$ is different for payload and preamble. FEC and interleaving are only applied to the payload.

(Ishibashi & Sugiura, 2014; da Rocha et al., 2017), but by doing this, SM actually loses one of its key advantages.

SpaRSe, however, can actually facilitate operation with a single transmit chain and practical pulse shaping filters. By simply inserting a number of zeros corresponding to the length of the pulse shaping filter after each spreading sequence, symbols from different antennas are completely separated in time. The rate loss that incurs from this is negligible, even for modest spreading factors such as 256.

3. Testbed Implementation

This section focuses on the receiver design and algorithms. Where possible, components and blocks from the IEEE 802.15.4 LECIM DSSS implementation in `gr-lpwan` (Wunsch & Maier, 2018) are reused. As the FD-DSM demodulator is implemented as a kernel, the basic structure can also easily be adapted to other modulation schemes. A block diagram of the receiver, which will be discussed in the remainder of this section, is given in Fig. 3.

3.1. Synchronization

The most computationally demanding part of the receiver is the synchronization stage. Incoming signals are usually far weaker than the thermal noise level, therefore requiring immense processing gain for detection, which also rules out simple techniques such as energy detection. Basically, a brute force approach is employed that considers every sample as a possible frame start and tries to demodulate the preamble. Furthermore, different frequency offset hypotheses may have to be checked at the base station due to low quality crystals in the low-cost sensors. The frequency spacing between the different hypotheses can be calculated as presented in (Wunsch et al., 2017).

Testing of multiple frequency hypotheses, however, is not yet implemented in the proof-of-concept testbed, where a manual presynchronization or a shared clock is required to guarantee that the frequency offset is within an acceptable range.

The synchronization stage, which operates at a sampling rate of $f_s = K \cdot R_c$ with K being the oversampling factor and R_c the chip rate, respectively, first performs the correlation of the input signal $\boldsymbol{r}$ with the real-valued and pulse-shaped preamble spreading sequence $\boldsymbol{c}_\mathrm{p}$ of length $L_{\boldsymbol{c}} = K \cdot \mathrm{SF} + N - 1$ with SF and N denoting the spreading factor and the length of the pulse shaping filter, respectively:

$$y_i = \sum_{n=0}^{L_\mathrm{c}-1} r_{n+i} c_{\mathrm{p},n} \tag{6}$$

If a preamble is present, this yields sharp correlation spikes corresponding to the preamble symbols separated by $L_{\boldsymbol{c}}$ samples. The correlator, which can also be implemented very efficiently using fast convolution, is followed by a parallel processing structure that divides the input signal in its polyphase components. There are $N_\mathrm{poly} = 2 \cdot L_{\boldsymbol{c}}$ branches which corresponds to the number of samples per STBC symbol.

The polyphase components on the n-th branch represent a subsampling of $\mathbf{y}$ by a factor of N_poly and can be written as $y_i^{(n)} = y_{n+i \cdot N_\mathrm{poly}}$. A decision on a possible frame start is made by continuously demodulating soft bits on each branch and correlating these with the preamble $\boldsymbol{p}$:

$$\rho_i^{(n)} = \sum_{k=0}^{L_{\boldsymbol{p}}-1} \mathrm{LLR}_{i+k}^{(n)} \cdot p_k \tag{7}$$

$$p_k \in \{-1, 1\},\ 0 \le k \le L_{\boldsymbol{p}} - 1 \tag{8}$$

For the sake of simplicity, the implementation assumes σ_n^2 to be constant over the entire frame and therefore neglects it

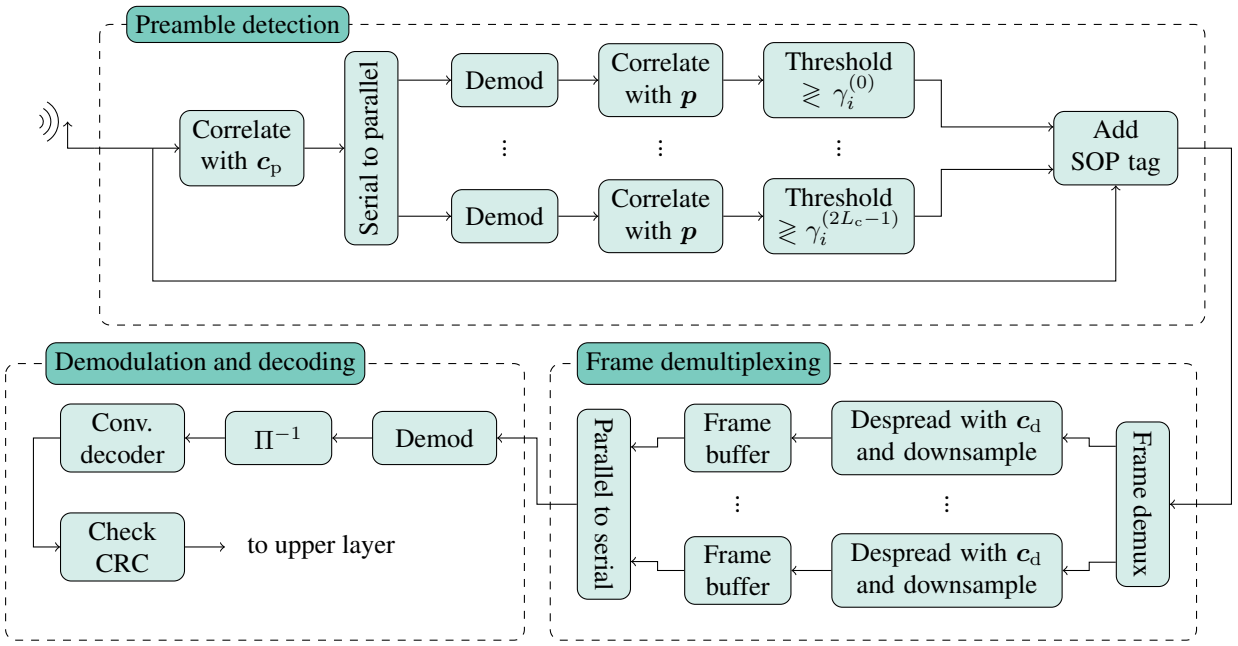

Figure 2. SpaRSe receiver block diagram. The preamble detector considers every sample as a possible start-of-frame and tries to demodulate the preamble, tagging the respective sample if the adaptive threshold is exceeded. Afterwards, overlapping frames are separated, despread, and the resulting symbols are buffered until the entire frame has been received. Contiguous frames are then demodulated, deinterleaved, decoded, and passed to the higher layer, if the final CRC check is successful.

for the calculation of the LLRs, also alleviating the requirement to estimate the Signal-to-Noise Ratio (SNR). Hence, the performance potentially degrades for non-stationary environments.

In the next step, the $\rho_i^{(n)}$ are compared to an adaptive threshold $\gamma^{(n)}$ which is calculated as

$$\gamma_i^{(n)} = \mu_{|\rho|,i}^{(n)} + \beta \cdot \sigma_{|\rho|,i}^{(n)} \text{ with} \tag{9}$$
$$\mu_{|\rho|,i}^{(n)} = (1-\alpha)\mu_{|\rho|,i-1}^{(n)} + \alpha|\rho_i^{(n)}|,$$
$$\sigma_{|\rho|,i}^{(n)} = \left((1-\alpha)(\sigma_{|\rho|,i-1}^{(n)})^2 + \alpha(|\rho_i^{(n)}| - \mu_{|\rho|,i}^{(n)})^2\right)^{\frac{1}{2}}$$

where β controls the false alarm probability. It has been verified in Monte Carlo simulations that the random variable which is created by averaging $|\rho_i^{(n)}|$ can be accurately modeled by a Normal distribution due to the Central Limit Theorem (CLT) if there is no preamble present. We use the absolute value of $\rho_i^{(n)}$ because the Autocorrelation Function (ACF) of $\boldsymbol{p}$ can take negative values and therefore would decrease the threshold even in the presence of a signal. It shall be noted that the ACF of $\boldsymbol{p}$ still exhibits considerable peaks at non-zero lags as shown in Fig. 3, easily leading to false alarms.

Once the threshold is exceeded, timing synchronization is achieved and the beginning of the frame is marked with a stream tag that also contains the last preamble symbol to allow error-free differential payload demodulation beginning

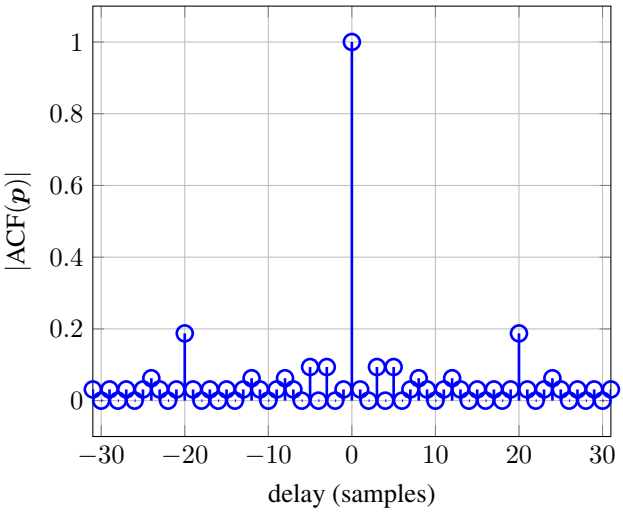

Figure 3. Magnitude of the normalized ACF of the preamble sequence $\boldsymbol{p}$ as defined in IEEE 802.15.4 for the LECIM DSSS PHY. The peaks at non-zero lags can cause a considerable false alarm rate, especially in high-SNR regimes.

Table 1. Measurement setup configuration

parameter	value
center frequency	$2.5\,\mathrm{GHz}$
chip rate	$1\,\mathrm{Mchip/s}$
pulse shaping	RRC with $\alpha_{\mathrm{RRC}} = 1$
sampling rate	$2\,\mathrm{MS/s}$ (TX), $4\,\mathrm{MS/s}$ (RX)
spreading factor	$\mathrm{SF} = 256$
payload length	$16\,\mathrm{byte}$ incl. CRC32
preamble length	$4\,\mathrm{byte}$
thresholding parameter β	off (perfect sync.), 3, 5
FD-DSM parameters	$L = 2$
frame duration	$76.6\,\mathrm{ms}$
data rate	$1.566\,\mathrm{kbit/s}$

at the first payload symbol. An example of the graphical output at the receiver is given in Fig. 4.

3.2. Payload Demodulation and Decoding

Before the payload can be demodulated, the symbols need to be extracted from the sample stream. In a multiple access scenario, it is also possible that frames from different users overlap. The demultiplexer performs the despreading with the payload spreading sequence $\boldsymbol{c}_{\mathrm{d}}$ and subsequent downsampling based on the frame start position marked by the tags from the preamble detector and passes only contiguous frames to the demodulator, thereby resolving potential overlaps.

After deinterleaving, the soft bits from the demodulator are then passed through the soft-input Viterbi decoder. Once decoded, the payload is checked for remaining bit errors with the trailing 4 byte CRC word. If the check passes, the CRC is stripped and the payload can be passed to the higher layers.

4. Measurements

To validate the proposed receiver design, over-the-air measurements were performed and compared to simulation results. The configuration of the transmitter and receiver is shown in Table 1.

4.1. Setup and Calibration

The measurement setup consists of a single laptop, two Universal Software Radio Peripheral (USRP) B210 devices, and an Octoclock-G (ett) from Ettus Research as reference source to avoid large frequency offsets. There is an unobstructed line-of-sight between the radios which are positioned about $2.5\,\mathrm{m}$ apart. Both radios use commercial, omni-directional Wi-Fi antennas. The receiving USRP employs an additional $40\,\mathrm{dB}$ attenuator while the antenna ports of the transmitting devices are attenuated by $10\,\mathrm{dB}$ to reduce the impact of possible interferers and to realize very low SNR settings while staying in the linear range of the amplifiers. The measurements were performed in an empty meeting room, i.e., no rapid channel fluctuations are expected. Even though the transmissions were recorded and demodulated offline to ensure reproducibility, it is worth mentioning that the system is real-time capable in the given configuration.

For calibration, the transmitter is set to a high gain setting ($80\,\mathrm{dB}$), where the influence of the additive noise power P_{n} on the signal power P_{s} at the receiver can be neglected, hence $P_{\mathrm{RX,s}} = P_{\mathrm{s}} + P_{\mathrm{n}} \approx P_{\mathrm{s}}$. $P_{\mathrm{RX,s}}$ denotes the received signal power after the input filter which is matched to the chips' pulse shape: an RRC pulse with roll-off factor $\alpha_{\mathrm{RRC}} = 1$. Due to the cyclostationarity of the modulated signal, however, the SNR at the demodulator (after reduction of the sample rate to the actual information symbol rate) is slightly higher than the average power at the output of the matched filter. For an oversampling factor of $K = 4$ this amounts to about $1.5\,\mathrm{dB}$ and can thus be taken into account for the SNR calculation. Turning off the transmitter, the remaining noise power $P_{\mathrm{RX,n}}$ can be used to calculate the SNR at the high-gain setting. Given the high linearity of the device in the gain region of interest, the SNR at the demodulator can then be controlled by adjusting the gain G accordingly using the following formula:

$$\mathrm{SNR} = G - 80\,\mathrm{dB} + \mathrm{SNR}_{80\,\mathrm{dB}} + 10\log_{10}(\mathrm{SF}) \quad (10)$$

with

$$\mathrm{SNR}_{80\,\mathrm{dB}} = 10\log_{10}\left(\frac{P_{\mathrm{RX,signal}}}{P_{\mathrm{RX,noise}}}\right) + 1.5\,\mathrm{dB}. \quad (11)$$

Taking into account the FEC rate $R = \frac{1}{2}$, we finally arrive at

$$\frac{E_{\mathrm{b}}}{N_0} = \mathrm{SNR} + 3.01\,\mathrm{dB}. \quad (12)$$

4.2. Results

Fig. 5 shows the measurement results for the described setup as well as a simulation with an all-ones channel (i.e., all-to-all and maximally correlated) and AWGN. A simulation with perfect synchronization serves as the baseline for the synchronization performance, which is evaluated with different values for β, i.e., different thresholds. The spreading factor is set to $\mathrm{SF} = 256$.

It can clearly be seen from the simulation results that the STBC can be decoded even for completely correlated channel coefficients. Furthermore, the measurement results are in good agreement with the simulated curves, hence validating the proposed synchronization algorithm. It should

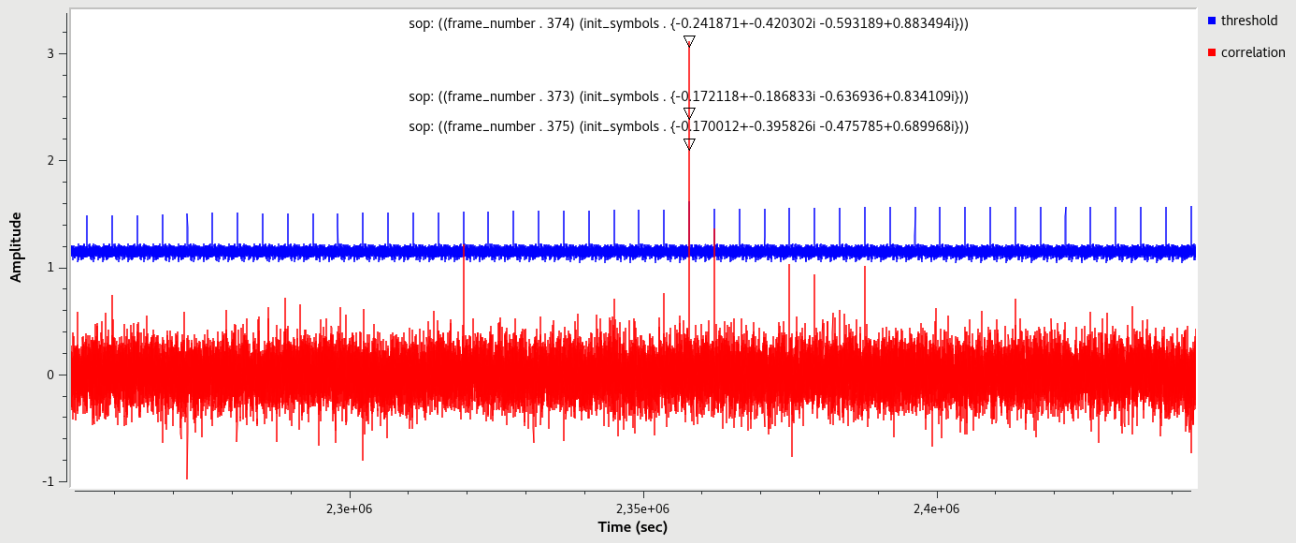

Figure 4. Visualization of the preamble detection process in the receiver for an over-the-air transmission. SF $= 256$, $\beta = 5$, $E_\mathrm{b}/N_0 \approx 7\,\mathrm{dB}$, i.e., the PER is on the order of 1%. It can be seen that the oversampling leads to multiple detections of the same packet. The thresholds for the different polyphase components are calculated independently which yields a higher threshold for the components that align with the detected packet due to the contribution of the ACF of the preamble $\boldsymbol{p}$, which can also be seen to effectively avoid some false alarms close to the main peak.

be noted that the simulation applies a simple AWGN channel while the system probably experienced fading with Rician characteristics and a very dominant line-of-sight component during the measurements; therefore, some deviation is to be expected.

Moreover, reducing the number of false alarms by setting the threshold relatively high with $\beta = 5$ degrades the performance only by an negligible amount but reduces the computational load considerably as it suppresses most of the peaks caused by the imperfect ACF.

5. Conclusion and Outlook

In this paper, we have presented our Software Defined Radio based testbed for *SpaRSe*, which is able to exploit transmit antenna diversity with a single transmit chain. We discussed its implementation, focusing on the receiver and especially on the synchronization stage that employs a polyphase approach with adaptive thresholding.

Furthermore, we provided the results of an indoor over-the-air verification test showing good agreement with simulation results, which are also provided. The synchronization process exhibits low implementation loss and effectively avoids false alarms.

At the time of writing, the receiver does not yet consider multiple frequency hypotheses as well as clock drift, which may arise due to low quality oscillators. Extending the

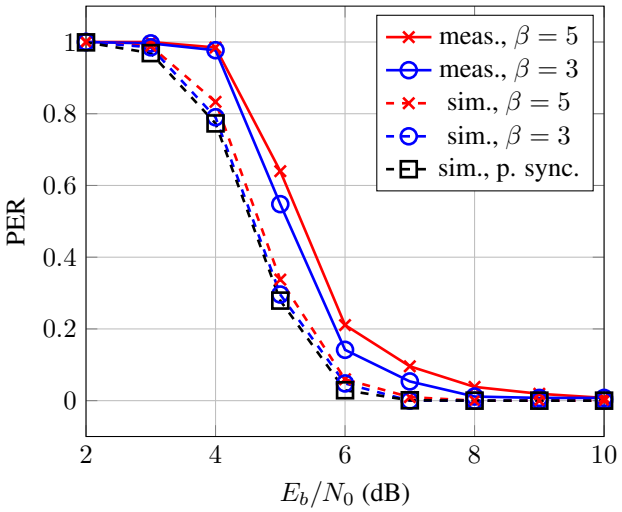

Figure 5. Comparison of over-the-air measurements and simulation results for the PER and SF $= 256$. Simulations and measurements show the same trend where lower β exhibit a performance marginally closer to the optimal case with perfect synchronization but also cause a considerably higher false alarm rate and therefore increased computational load.

synchronization stage with a modulated filter bank will therefore be the subject of future work. As the majority of the computational load in the receiver is caused by the synchronization stage and the structure is inherently parallel, offloading the work to multiple cores or a Field-Programmable Gate Array would also bring a significant increase in throughput. Further improvements will comprise a Delay-Locked Loop for the tracking of clock drift for very high spreading factors and the implementation of an incoherent RAKE receiver structure to combine delayed signals in case of frequency-selective fading.

References

Ettus research product overview. https://www.ettus.com/product. Online, accessed 2018/18/06.

RPMA white paper. https://www.ingenu.com/portfolio/how-rpma-works-white-paper/. Accessed: 2018-27-06.

da Rocha, C. A. Faria, Ucha-Filho, B. F., and Ruyet, D. Le. Study of the impact of pulse shaping on the performance of spatial modulation. In *2017 International Symposium on Wireless Communication Systems (ISWCS)*, pp. 303–307, Aug 2017. doi: 10.1109/ISWCS.2017.8108129.

Ishibashi, K. and Sugiura, S. Effects of Antenna Switching on Band-Limited Spatial Modulation. *IEEE Wireless Communications Letters*, 3(4):345–348, Aug 2014. ISSN 2162-2337. doi: 10.1109/LWC.2014.2315819.

Jeganathan, J., Ghrayeb, A., and Szczecinski, L. Spatial modulation: optimal detection and performance analysis. *IEEE Communications Letters*, 12(8):545–547, Aug 2008. ISSN 1089-7798. doi: 10.1109/LCOMM.2008.080739.

Martin, P. A. Differential Spatial Modulation for APSK in Time-Varying Fading Channels. *IEEE Communications Letters*, 19(7):1261–1264, July 2015. ISSN 1089-7798. doi: 10.1109/LCOMM.2015.2426172.

Mesleh, R. Y., Haas, H., Sinanovic, S., Ahn, C. W., and Yun, S. Spatial Modulation. *IEEE Transactions on Vehicular Technology*, 57(4):2228–2241, July 2008. ISSN 0018-9545. doi: 10.1109/TVT.2007.912136.

Raza, U., Kulkarni, P., and Sooriyabandara, M. Low Power Wide Area Networks: An Overview. *IEEE Communications Surveys Tutorials*, 19(2):855–873, Secondquarter 2017. doi: 10.1109/COMST.2017.2652320.

Wunsch, Felix and Maier, Kristian. Source code repository. https://github.com/kit-cel/gr-lpwan, 2018.

Wunsch, Felix, Maier, Kristian, Jäkel, Holger, and Jondral, Friedrich K. Implementation and Performance Evaluation of IEEE 802.15.4 LECIM DSSS PHY at 2.4 GHz. *Proceedings of the GNU Radio Conference*, 2 (1), 2017. URL https://pubs.gnuradio.org/index.php/grcon/article/view/19.

Wunsch, Felix, Jäkel, Holger, and Jondral, Friedrich K. A Low-Complexity Air Interface With Transmit Diversity for Low Power Wide Area Networks. *VTC Spring 2018 (Porto)*, accepted for publication, 2018a.

Wunsch, Felix, Ströer, Max, Jäkel, Holger, and Jondral, Friedrich K. Low Power Wide Area Networks at 2.4 GHz: A Viable Option? *VTC Fall 2018 (Chicago)*, accepted for publication, 2018b.

Zhang, W., Yin, Q., and Deng, H. Differential Full Diversity Spatial Modulation and Its Performance Analysis With Two Transmit Antennas. *IEEE Communications Letters*, 19(4):677–680, April 2015. ISSN 1089-7798. doi: 10.1109/LCOMM.2015.2403859.