

On the Reception of Concurrent Transmissions in Wireless Sensor Networks

Matthias Wilhelm, *Student Member, IEEE*, Vincent Lenders, *Member, IEEE*, and Jens B. Schmitt

Abstract—Numerous studies have shown that concurrent transmissions can help boost wireless network performance despite the possibility of packet collisions. However, while these works provide empirical evidence that concurrent transmissions may be received reliably, existing signal capture models only partially explain the root causes of this phenomenon. We present a comprehensive mathematical model for MSK-modulated signals that makes the reasons explicit and thus provides fundamental insights into the key parameters governing the successful reception of colliding transmissions. A major contribution is the closed-form derivation of the receiver bit decision variable for an arbitrary number of colliding signals and constellations of power ratios, time offsets, and carrier phase offsets. We systematically explore the factors for successful packet delivery under concurrent transmissions across the whole parameter space of the model. We confirm the capture threshold behavior observed in previous studies but also reveal new insights relevant to the design of optimal protocols. We identify capture zones depending not only on the signal power ratio but also on time and phase offsets.

Index Terms— Packet collisions, MSK modulation, co-channel interference, capture zone, Glossy.

I. INTRODUCTION

CONVENTIONAL wireless communication systems consider packet collisions as problematic and try to avoid them by using techniques like carrier sense, channel reservations (virtual carrier sense, RTS/CTS handshakes), or arbitrated medium access (TDMA, polling). The intuition is that concurrent transmissions cause irreparable bit errors at the receiver and render packet transmissions undecodable. However, researchers have found that this notion is too conservative. If the power of the signal of interest exceeds the sum of interference from colliding packets by a certain threshold, packets can in general still be received successfully despite collisions at the receiver. This effect, referred to as the *capture effect* [1], has been explored extensively and validated in many independent practical studies on various communication systems such as IEEE 802.11 [2]–[5] and IEEE 802.15.4 [6]–[8].

Manuscript received November 28, 2013; revised April 24, 2014; accepted August 5, 2014. Date of publication August 20, 2014; date of current version December 8, 2014. The associate editor coordinating the review of this paper and approving it for publication was S. Valaee.

M. Wilhelm and J. B. Schmitt are with the Department of Computer Science, University of Kaiserslautern, 67663 Kaiserslautern, Germany (e-mail: wilhelm@cs.uni-kl.de; jschmitt@cs.uni-kl.de).

V. Lenders is with armasuisse Science and Technology, 3602 Thun, Switzerland (e-mail: vincent.lenders@armasuisse.ch).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TWC.2014.2349896

Over the past years, the view on packet collisions has therefore changed considerably. Since it is possible for some or even all packets in a collision to survive, there are opportunities to increase the overall channel utilization and to improve the network throughput by designing protocols that carefully select terminals for transmitting at the same time [9], [10]. The benefits and potential performance improvements of concurrent transmission are not just of theoretical interest but have been demonstrated practically and adopted in application areas such as any-cast [11], [12], neighbor counting [13], or rapid network flooding [14]–[18], especially in the context of wireless sensor networks (WSNs).

Although protocols that exploit concurrent transmissions have shown the potential to boost the overall performance of existing wireless communication systems, their success cannot be explained with capture threshold models based on the Signal to Interference and Noise Ratio (SINR) alone. Recent studies have shown that, while the relative signal powers of colliding packets indeed play an important role in the reception probability, other factors are also of major importance. For example, several experimental studies report that the relative timing between colliding packets has a significant influence on the reception probability [5], [19]. Others report that the coding [20] or packet content [11] may also greatly influence the reception performance in the presence of collisions. Further factors such as the carrier phase offset between a packet of interest and colliding packets also need to be considered [21].

In this paper, we strive to provide a comprehensive model accounting for all these factors, focusing on packet collisions in IEEE 802.15.4 based WSNs. Such a model will allow protocol designers to better understand the root causes of packet reception and exact conditions under which concurrent transmissions actually work, and thus to design optimal protocols based on these factors. While previous studies [3], [7], [22]–[24] also looked at factors that determine the success of concurrent packet reception, these works are either based on practical experiments and have therefore led to empirical models that cannot be generalized easily, or derived simplified models that do not account for all impact factors. This work advances the field by providing a unified analytical model accounting for the major factors identified above (see also Section II). Our model (\rightarrow Section III) is based on a mathematical representation of the physical layer using continuous-time expressions of the IQ signals entering the receiver's radio interface. This fundamental and comprehensive model allows to represent an arbitrary number of colliding packets as a linear superposition of the incoming signals.

A major contribution of this work is a closed-form analytical representation of the bit decision variable at an optimal receiver's demodulator output based on these IQ signals (\rightarrow Section IV). This result enables the deterministic computation of the bit demodulation decision and hence to compute the actual performance of concurrent transmissions for any colliding parameter constellations. Having a bit-level model of reception is not only beneficial for the comprehension of the collision process, it also contributes to application areas where a precise bit-level analysis is needed, such as partial packet reception [25], understanding bit error patterns in low-power wireless networks [26], [27], or signal manipulation attacks at the physical layer [21].

Using our model, we explore the parameter space of the reception of MSK-modulated colliding packets considering both uncoded and Direct Sequence Spread Spectrum (DSSS) based systems (\rightarrow Section V), analyzing the influence of the parameters on the resulting packet reception ratio (PRR) for concurrent transmissions. While the analysis shows that our model agrees with experimental results in the literature, it also provides much more detailed insights into the performance characteristics of protocols that exploit collisions [11], [14]–[18]. In particular, we show that the good performance of these protocols should be attributed equally to coding (e.g., DSSS) and power capture. In addition, based on our analysis we identified parameter constellations where concurrent transmissions work reliably. We therefore propose a generalization of the traditional capture threshold model based on the power ratios towards a *capture zone*. Capture zones result from the model insight that reception success does not depend on the power ratio between interfering signals alone, but on the time and phase offsets of sender and receiver as well.

To show the validity and accuracy of our model, we implemented and experimented with an application that is strongly dependent on physical layer characteristics, the reception of unsynchronized signals. We performed this experiment with two widely used commercial IEEE 802.15.4 receiver implementations (TI CC2420 and Atmel AT86RF230) to demonstrate that our results are receiver-independent (\rightarrow Section VI). The results validate our claim that our model accurately captures the behavior of realistic receivers in the face of concurrent transmissions. Finally, we discuss parameter settings for an optimal protocol design (\rightarrow Section VII).

II. IMPACT FACTORS

Different factors influence the probability of a successful reception under collisions. This section discusses the main factors that have been identified in the literature. Subsequently, we consider them jointly in our mathematical model to predict the outcome of concurrent transmissions.

Power Ratio: The signal power is a crucial factor for successful reception in general, and it plays a major role in the reception under collisions as well. SINR-based models are widely used to model the packet reception in a shared medium, for example in the Physical Model [28] and its variants [29], [30]. The classical SINR model states that a stronger signal

is received if its signal power P_s exceeds the channel noise P_n and the sum of interfering signal powers $\sum_i P_i$ by a given threshold, i.e.,

$$\frac{P_s}{P_n + \sum_i P_i} > \delta_{\text{SINR}}.$$

This simple model is accurate for uncorrelated interfering signals such as additive white Gaussian noise (AWGN). However, when the interference is correlated (such as colliding packets), this model is not always accurate and further factors must be considered [3], [5], [19].

Signal timing: The relative timing of colliding packets greatly influences the reception process. This is because the receiver locks onto a packet during the synchronization phase at the start of the transmission. If a stronger signal arrives later, it captures the receiver and disturbs the first packet reception, and both packets in the collision are lost. Thus, in packet radios, power capture alone is not sufficient for successful reception, rather the receiver must be synchronized and locked onto the captured signal as well. Several research contributions analyze possible collision constellations and their effect on packet reception [5], [19], and propose a new receiver design that releases the lock when a stronger packet arrives, discards the first and receives the second packet, the so-called *message-in-message (MIM) capture* [5], [22]. Subsequent works apply these insights to improve network throughput. For example, Manweiler *et al.* [31] propose collision scheduling to ensure that MIM is leveraged, thus increasing spatial reuse.

Channel coding: A further factor that influences packet reception success is bit-level coding. For example, in DSSS systems a group of b bits is encoded into a longer sequence of B chips [32]. The benefit of this approach is that resilience to interference is increased because the chipping sequences can be cross-correlated at the receiver, which effectively filters out uncoded noise. However, DSSS systems require interfering signals to be uncorrelated, e.g., signals without coding or with orthogonal chipping sequences (as in CDMA), to achieve their theoretical coding gain. Another possibility is a sufficient time offset between interfering packets with the same coding; this phenomenon is known as *delay capture* [20]. As networking standards such as IEEE 802.11 and IEEE 802.15.4 generally use DSSS with identical codes for all participants, existing experimental works on collisions and capture observe the effects of DSSS implicitly.

Packet contents: Experimental results show that packets with identical payload and aligned starting times result in good reception performance and reduced latency in broadcast scenarios. For example, Dutta *et al.* [11] show that short packets can be received in such collisions with a PRR over 90%, thus enabling the design of an efficient receiver-initiated link layer. Similarly, the latency of flooding protocols widely used in WSNs can be greatly reduced [15], [17]. In these works, experiments in IEEE 802.15.4 networks reveal that the tolerable time offset between concurrent messages is small (approx. 500 ns), which adds challenges to protocol design and implementation. These insights also show that capture and packet synchronization alone are not sufficient

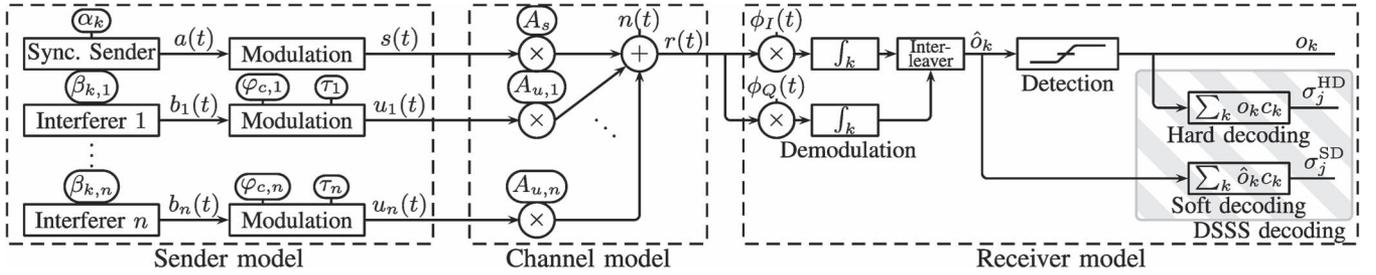


Fig. 1. System model, its parameters are shown in ovals (information bits α_k, β_k , carrier phase offset φ_c , time offset τ , signal amplitudes A_s, A_u). We consider one synchronized sender and n interferers on a collision channel that is the input to a receiver. Here, three channel coding schemes are considered, (i) uncoded, (ii) DSSS with hard decision decoding (HDD), and (iii) DSSS with soft decision decoding (SDD); resulting in different receiver paths.

to explain the performance of these protocols, and bit-level modeling that also includes signal timing and content is necessary.

Carrier phase: Considering the reception of bits at the physical layer, knowledge of the carrier phase at the receiver is crucial for successful reception of phase modulated signals because the information is carried in the phase variations of the signal, such that these offsets should be minimized [32]. Typically this is achieved during the synchronization phase of packet reception, and thus existing capture models have omitted phase offsets. However, there are two reasons why this is not sufficient. First, in novel protocols exploiting packet collisions, the synchronization during the preamble is not always able to succeed. Second, there are other new applications of concurrent transmissions that try to abandon the synchronization procedure. For example, Pöpper *et al.* [21] investigate the possibility of manipulating individual message bits on the physical layer, and conclude that carrier phase offsets are the major hindrance to do so reliably.

III. SYSTEM MODEL

In this section, we discuss the system model underlying our analysis, as shown in Fig. 1. It considers all factors from the previous section. From a bird's eye view, the model consists of three components: (i) the sender model that modulates the physical layer signals of $n + 1$ transmitters, one fully synchronized signal of interest (SoI) and n interferers with possibly differing transmission starting times and payloads; (ii) the channel model with all senders sharing a single collision channel that outputs a scaled superposition of all signals (according to their corresponding power at the receiver), and (iii) the receiver model with three detection methods: uncoded, DSSS with hard decision decoding (HDD), and DSSS with soft decision decoding (SDD). In the following, we discuss each component in detail.

A. Sender Model

In the first component, we modulate the physical signals of $n + 1$ senders. We instantiate our model with the Minimum Shift Keying (MSK) modulation, a widely used digital modulation with desirable properties, and of special interest because of its use in the 2.4 GHz PHY of IEEE 802.15.4 [33, Sec. 6.5], but

we also discuss other modulation schemes including O-QPSK, QPSK, and BPSK. For the signal representations, we follow the notation of Proakis and Salehi [32, Sec. 4.3].

1) *Synchronized Sender:* We assume that the receiver is fully synchronized to the SoI, i.e., the synchronization process has successfully acquired this signal and all interferers have relative offsets to it. The signal is then given by

$$s(t) = a_I(t) \cos\left(\frac{\pi t}{2T}\right) \cos \omega_c t + a_Q(t) \sin\left(\frac{\pi t}{2T}\right) \sin \omega_c t. \quad (5)$$

The signal consists of two components, the in- (I) and the quadrature-phase (Q) components. Modulated onto each component are the information signals (carrying the bits represented by $\alpha_k^I, \alpha_k^Q \in \{\pm 1\}$) given by

$$a_I(t) = \sum_{k=-\infty}^{\infty} \alpha_k^I \Pi\left(\frac{t - 2kT}{2T}\right) \quad (6)$$

$$a_Q(t) = \sum_{k=-\infty}^{\infty} \alpha_k^Q \Pi\left(\frac{t - (2k + 1)T}{2T}\right), \quad (7)$$

which represents a train of unit pulses Π with duration $2T$, the bit duration of the modulation (e.g., $2T = 1 \mu\text{s}$ in IEEE 802.15.4). The information signals are staggered, i.e., the Q-phase information signal is delayed by T in $a_Q(t)$. These signals are then shaped with half-sine pulses of duration $2T$, and modulated onto a carrier with frequency $\omega_c/2\pi$ (e.g., 2.4–2.48 GHz in IEEE 802.15.4). In the following, we use the angular frequency of baseband pulses $\omega_p = \pi/2T$, such that the first cosine term in (5) may be represented by $\cos \omega_p t$.

2) *(Unsynchronized) Interferers:* In addition to the synchronized sender, we consider n interferers transmitting concurrently, using the same modulation. These signals may not be synchronized to the receiver and each may carry its own payload. This introduces three additional parameters that influence the signal, the time offset τ_i , the carrier phase offset $\varphi_{c,i}$, and the information bits $\beta_{k,i}$. With a positive τ_i , an interfering signal arrives later at the receiver than the synchronized signal. The signal at the receiver for interferer i is given by

$$u_i(t; \tau_i, \varphi_{c,i}) = b_{I,i}(t - \tau_i) \cos \omega_p(t - \tau_i) \cos(\omega_c t + \varphi_{c,i}) \\ + b_{Q,i}(t - \tau_i) \sin \omega_p(t - \tau_i) \sin(\omega_c t + \varphi_{c,i}).$$

We assume that the phase offsets $\varphi_{c,i}$ are constant for the duration of a packet, i.e., there is no carrier frequency offset during a transmission. In our experiments in Section VI, we show that this assumption is reasonable because receiver implementations are compensating for possible drifts. For convenience, we express the pulse phase offset caused by τ_i as $\varphi_{p,i} = \omega_p \tau_i$.

3) *Other Modulation Schemes*: While our results are derived for the MSK modulation, it is possible to adapt them to other variants of the phase shift keying (PSK) modulation. We briefly describe the differences to major variants and highlight how these affect the analysis. Further details on the relationship between PSK modulation schemes can be found in Proakis and Salehi [32] and Pasupathy [34].

Offset QPSK: O-QPSK with a half-sine pulse shape is identical to MSK [35] and the results therefore also apply for this modulation. If O-QPSK is used in combination with rectangular pulse shaping instead, the signal is then given by

$$s_{\text{O-QPSK}}(t) = \frac{1}{\sqrt{2}} (a_I(t) \cos \omega_c t + a_Q(t) \sin \omega_c t).$$

The altered pulse shape leads to the omission of the factor $\cos \omega_p t$ present in (5), because the rectangular shaping is already included in the information signal $a(t)$. This leads to a simplification of our MSK results because pulse phase offsets φ_p that are caused by the time offset τ are not present.

Quadrature PSK: Considering QPSK, the change from O-QPSK is the missing time shift T in the quadrature phase. This leads to a different information signal for the Q phase,

$$a'_Q(t) = \sum_{k=-\infty}^{\infty} \alpha_k^Q \Pi \left(\frac{t - 2kT}{2T} \right).$$

When adapting our results to QPSK, this affects the indices k of the colliding bits.

Binary PSK: This scheme considers only the in-phase components of QPSK, its signal is given by

$$s_{\text{BPSK}}(t) = \frac{1}{\sqrt{2}} a_I(t) \cos \omega_c t.$$

This simplifies the derivations and results further, because there is no contribution from the Q phase signal in collisions.

B. Channel Model

In our model, we use an additive collision channel. The relation for the output signal is

$$r(t) = A_s s(t) + \sum_{i=1}^n A_{u,i} u_i(t; \tau_i, \varphi_{c,i}) + n(t).$$

Each signal is scaled by a positive, real-valued factor A , which contains both, possible signal amplifications by the sender and path loss effects that reduce the power at the receiver. In our evaluation, we use the Signal to Interference Ratio (SIR) at the receiver, given by $\text{SIR} = A_s^2 / (\sum_{i=1}^n A_{u,i}^2)$, to characterize the power relationship of the interfering signals. The contribution of all noise effects is accumulated in the linear noise term

$n(t)$; possible instantiations are a noiseless channel or a white Gaussian noise channel.

C. Receiver Model

In the final component of the model, we feed the signals' superposition $r(t)$ into an optimal receiver to discern the detected bits. The signal is demodulated and fed into one of three detector implementations: one for uncoded bits, and two variants of DSSS decoding.

1) *Demodulation*: Demodulation is performed for I and Q individually and the bits are then interleaved. We limit our discussion to the I component for brevity.

We use the matched filter function $\phi_I(t) = (2/T) \cos \omega_p t \cos \omega_c t$ and low-pass filtering for downconversion and demodulation, which is the optimal receiver for noiseless and Gaussian channels in the sense that it minimizes the bit error probability [32, Sec. 4.3]. The received signal $r(t)$ is multiplied by $\phi_I(t)$ and integrated for each bit period k to form the decision variable

$$\hat{o}_k^I = \Lambda_r^I(k) = \int_{(2k-1)T}^{(2k+1)T} r(t) \phi_I(t) dt.$$

The resulting (real) value is called *soft bit*. Because the combination of the interferers in the received signal is linear, the individual contributions can be divided into integrals for each signal:

$$\hat{o}_k^I = \Lambda_s^I(k) + \sum_{i=1}^n \Lambda_{u_i}^I(k) + \Lambda_n^I(k).$$

In our analytical evaluation in the following section, we derive closed-form expressions for $\Lambda_{u_i}^I$ and $\Lambda_{u_i}^Q$ to analyze the receiver output after a signal collision.

We point out that this simplified model does not include receiver-side techniques such as Automatic Gain Control (AGC) or phase tracking; however, we conjecture that the reception performance is still comparable. In fact, as our experiments in Section VI show, this assumption is justified and the simplified model is able to predict the reception behavior of real-world receiver implementations with good accuracy. We leave the investigation on the effects of these advanced techniques to future work.

2) *Uncoded Bit Detection*: The detection operation for uncoded transmissions is slicing, essentially a sign operation on the demodulation output, which results in binary output $o_k \in \{\pm 1\}$. Thus, a bit of the SoI is flipped if the contribution of the interferers changes the bit's sign.

3) *DSSS Decoding*: For coded transmissions, the number of chips exceeds the bits in a symbol, i.e., even if several chips are flipped it is still possible to decode a symbol correctly. We consider 2^b symbols ξ with chipping sequence c_ξ , each with a block length of B bit (i.e., the number of chips). For example, we have $b = 4$, $B = 32$ in IEEE 802.15.4.

We differentiate two modes of operation for the DSSS decoder, namely hard decision decoding (HDD) and soft decision decoding (SDD) [32].

TABLE I
ANALYTICAL RESULTS

No offsets	$\Lambda_u^I(k) = A_u \beta_k^I$	(1)
Carrier phase offset φ_c	$\Lambda_u^I(k) = A_u \left(\cos \varphi_c \beta_k^I - \frac{1}{\pi} \sin \varphi_c \left(\beta_{k-1}^Q - \beta_k^Q \right) \right)$	(2)
Time offset τ	$\Lambda_u^I(k) = \frac{1}{2T} A_u \left(\cos \varphi_p \left(\tau \beta_{k'-1}^I + (2T - \tau) \beta_{k'}^I \right) - \frac{2T}{\pi} \sin \varphi_p \left(\beta_{k'-1}^I - \beta_{k'}^I \right) \right)$	(3)
Carrier phase + time offset	$\Lambda_u^I(k) = \frac{1}{2T} A_u \left\{ \cos \varphi_c \left(\cos \varphi_p \left(\tau \beta_{k'-1}^I + (2T - \tau) \beta_{k'}^I \right) - \frac{2T}{\pi} \sin \varphi_p \left(\beta_{k'-1}^I - \beta_{k'}^I \right) \right) \right. \\ \left. - \sin \varphi_c \left(\sin \varphi_p \left(\tau \beta_{k'Q'-1}^Q + (2T - \tau) \beta_{k'Q'}^Q \right) + \frac{2T}{\pi} \cos \varphi_p \left(\beta_{k'Q'-1}^Q - \beta_{k'Q'}^Q \right) \right) \right\}$	(4)

Contributions of an interfering signal to the demodulator output Λ_u^I for the k -th bit. The results present the relationship between in- and quadrature phase bits sent by an interferer (β^I and β^Q) with time-adjusted bit indices $k', k^{Q'}$; it also considers the effects of carrier phase offsets φ_c and time offsets (φ_p and τ). The corresponding notation is introduced in Sections III and IV.

Hard Decision Decoding: In *HDD*, the decoder uses sliced (binary) values o_k as its input, and then chooses the symbol with the highest bit-wise cross-correlation of all chipping sequences. In this way, *HDD* can be viewed as an additional step that takes a group of uncoded bits with B elements (from the uncoded bit detection described above) to determine a symbol σ_j^{HD} , i.e., a group of b bits. For *HDD*, the decoder is given by

$$\sigma_j^{\text{HD}} = \arg \max_{0 \leq \xi < 2^b} \left| \sum_{k=0}^{B-1} o_{jB+k} C_{\xi,k} \right|. \quad (8)$$

Soft Decision Decoding: In *SDD*, the real-valued, unquantized demodulator output \hat{o}_k (*soft bits*) is used as decoder input directly, in contrast to the binary values o_k used in *HDD*. This is beneficial because soft bits provide a measure of detection confidence and demodulation quality, and thus adds weighting to the bits used in the cross-correlation.

IV. MATHEMATICAL ANALYSIS

Based on the system model in Fig. 1, we analyze the contributions of each interfering signal to the overall demodulator output; the sum of these contributions is the decision variable of bit detection. We first present the general case considering all system parameters in Theorem 1. Subsequently, we illustrate its interpretation using selected parameter combinations.

Theorem 1: For an interfering MSK signal $u(t)$ with offset parameters τ and φ_c , the contribution to the demodulation output $\Lambda_u^I(k)$ is given by (4), shown in Table I.¹

Due to space restrictions, the proof of this theorem can be found in [36]. To provide a better understanding of the effects of the parameters, we focus on selected parameter constellations and discuss the resulting equations. Then we revisit Theorem 1 and discuss the combination of effects.

1) **Synchronized Signal:** In the simplest case both offsets, time and phase, are zero, i.e., the interfering signal is also fully synchronized to the receiver. The result is given in (1) in Table I. The signal's contribution to the k -th bit is $\Lambda_u^I(k) = A_u \beta_k^I$.

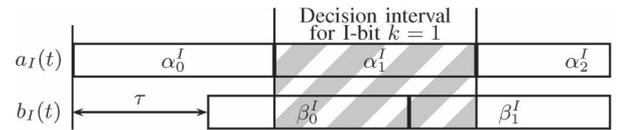


Fig. 2. Example of a time offset τ during detection: three bits influence the bit decision on the second bit in a collision between two signals. The active bits in the decision interval are highlighted (the synchronized sender's bit α_1^I and interferer's in-phase bits β_0^I and β_1^I).

The bit decision of bit k , i.e., the sign of the equation, is governed by β_k^I . The magnitude of the contribution is controlled by the amplitude of the signal A_u , and thus stronger signals lead to a greater contribution to the decision variable \hat{o}_k . As an example, consider two signals $s(t)$ and $u(t)$ that are both fully synchronized to the receiver. The detector output of bit k is then $A_s \alpha_k^I + A_u \beta_k^I$. If both senders transmit the same bit ($\alpha_k^I = \beta_k^I$), then the signals interfere constructively and push the decision variable further away from zero. If, on the other hand, the bits are different, then the decision variable has the sign of the stronger signal; this is the well-known power capture effect for a single bit.

2) **Carrier Phase Offset:** Next, we analyze the effect of carrier phase offsets when the signals are fully time-synchronized ($\tau = 0$). The result is given in (2), shown in Table I. We observe two effects of the carrier phase offset. First, the bit contribution of β_k^I is scaled by $\cos \varphi_c \leq 1$, which leads to reduced absolute values (and thus a smaller contribution to the decision variable) and potentially causes the bit β_k^I to flip for $\varphi_c \in ((\pi/2), (3\pi/2))$. Second, the quadrature phase starts to leak into the decision variable and thus two additional bits β_{k-1}^Q, β_k^Q influence the outcome. This contribution, however, is scaled by $\pi^{-1} \sin \varphi_c$, and only appears when the two Q bits are alternating during the integration interval. In essence, uncontrolled carrier phase offsets may lead to unpredictable bits in the detector output because of carrier phase offset induced bit flips.

3) **Time Offset:** If the signals are phase-matched but shifted in time, the detector output is given by (3), shown in Table I. We make three observations here. The bit index k needs to be adjusted because bits may be time-shifted into the integration interval, see Fig. 2; the new index is given by $k' = k - \lfloor \tau/2T \rfloor$, with $\lfloor \cdot \rfloor$ denoting the floor function. We call these *active bits*

¹We omit the subscript i for clarity in the equations. The results for the quadrature phase are given by the same equations when the roles of I and Q are exchanged.

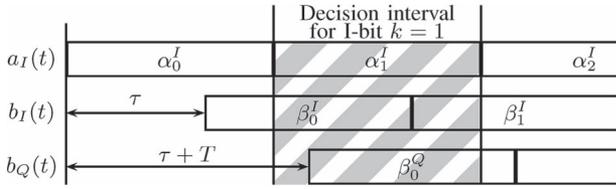


Fig. 3. Example of time and phase offsets combined: the decision on the second bit ($k = 1$) is influenced by four bits in this example (the synchronized sender's bit α_1^I , the interferer's in-phase bits β_0^I and β_1^I due to time shifts, and quadrature bit β_0^Q from carrier phase offsets).

because they contribute to the bit decision. These bits overlap partially or fully, and their active time duration is $\tau^Q = \tau - 2\lfloor \tau/2T \rfloor T$, the underscore signifies that its value is confined to the interval $[0; 2T)$. However, these bits do not contribute to the decision directly but are scaled by $\cos \varphi_p$, which is caused by the half-sine pulse shaping of MSK. This scaling means that bit contributions are diminished and may be flipped by certain time offsets. Finally, a term scaled by π^{-1} is introduced that is only present when bits are alternating. However, these bits are the same in-phase bits $\beta_{k'-1}^I, \beta_{k'}^I$, the Q phase does not leak in this setting.

4) *Both Offsets*: Finally, when both offsets are present as in Theorem 1, we can interpret the result as a combination of the above effects. A graphical illustration of the active bits is shown in Fig. 3. Due to the staggering of bits (the Q bits are delayed by T), the indices of leaking bits of the Q phase also need to be adjusted, the new index is $k^{Q'} = k - \lfloor (\tau + T)/2T \rfloor$, and the active time interval τ^Q is derived similarly to above.

In summary, we observe that the contribution of the interfering signal is complex and that φ_c and φ_p can potentially flip the original bits β_k^I . This should be bad news for collision-aware protocols that use identical payload to achieve constructive interference (e.g., SCIF [17]): these bits can flip easily and then generate destructive interference. However, coding helps to alleviate these negative effects as we will see in the next section.

V. PARAMETER SPACE EXPLORATION

Equipped with the closed-form analytical model of the bit-wise receiver outputs, we systematically explore the parameter space of the reception of concurrent transmissions in detail.

A. Methodology

To numerically study the transmission reception success under interference we perform so-called Monte Carlo simulations (see Jain [37]); that means we do time-static simulations of independent packet transmissions in which we randomly vary the analytical model's parameters to investigate their influence on performance parameters such as packet reception ratio, bit and symbol error rate. Conceptually, the simulator is just a software version of the mathematical model (written in Python) applied to a whole packet; it is not meant to validate the model but to experiment with randomly chosen values for the model parameters and to provide more insights on the success probability of concurrent transmissions. The simulation code

is available for download at <http://disco.cs.uni-kl.de/content/collisions>; there, the interested reader can also find an interactive visualization of the model.

For most experiments, the time offset between sender and interferer is fixed and is our primary factor in the numerical analysis, i.e., in the plots we show the reception performance depending on the time offset. The other parameters of the model are treated as secondary factors and are randomly varied. Generating 1000 independent packet transmissions for each data point in the presented graphs thus represents the secondary factors' average contribution to the reception success. We provide more details on the choices for the model's parameters for sender and channel in the following.

1) *Sender Model*: For ease of presentation, we mainly consider the presence of one synchronized sender and one interferer; we denote these parties as \mathcal{S} and \mathcal{I} with signals $s(t)$ and $u(t)$, respectively. In Section V-B3, we consider the n interferer case separately. We analyze the reception performance of groups of associated bits, or packets; in this case, a single bit error leads to a packet drop. The packet reception ratio (PRR) is the fraction of packets that arrive without errors divided by the total number of packets. We use packets with a length of 64 bit. We consider two categories of colliding packets, either with independent (\mathcal{S} and \mathcal{I} trying to exploit spatial reuse) or identical content ($\alpha_k = \beta_k$, as it is the case for collision-aware flooding protocols). The bits to send are chosen in the following manner: for uncoded transmissions, α_k is drawn bitwise i.i.d. from a Bernoulli distribution over $\{-1, 1\}$, and either the same procedure is performed for β_k (independent packets) or simply copied over from α_k (identical packets). For coded packets, we draw symbols i.i.d. uniform random from $\{0, \dots, 15\}$ and spread these symbols according to the chipping sequences defined by the IEEE 802.15.4 standard [33, Sec. 6.5]. This means that 4 bit groups are first spread to 32 bit chipping sequences before they are transmitted in α_k, β_k .

In accordance to the literature [38], as the carrier phase offset is hard to control because of oscillators drifts and other phase changes during transmission, we draw φ_c i.i.d. uniform randomly from $[0; 2\pi)$ for each packet unless stated otherwise. On the other hand, we use the same time offset τ for all packets because experimental work shows that this timing can be precisely controlled. For example, Glossy [15] achieves a timing precision of 500 ns over 8 hops with 96% probability, and Wang *et al.* [18] report a 95% percentile time synchronization error of at most 250 ns. For our simulations, we used 1000 packets for each value of τ .

2) *Channel Model*: To concentrate on the impact of signal interference, we consider a noiseless channel. This is a well-accepted assumption when both signals are significantly above noise floor level [39, Sec. 8]. We set $A_s = 1$ and $A_u = \text{SIR}^{-(1/2)}$.

B. Reception of the Synchronized Signal of Interest

1) *Capture Threshold Under Independent Payload*: In our first case study, we consider the transmission of independent payload. This situation occurs, e.g., when two uncoordinated senders detect a clear channel, transmit, and the packets collide

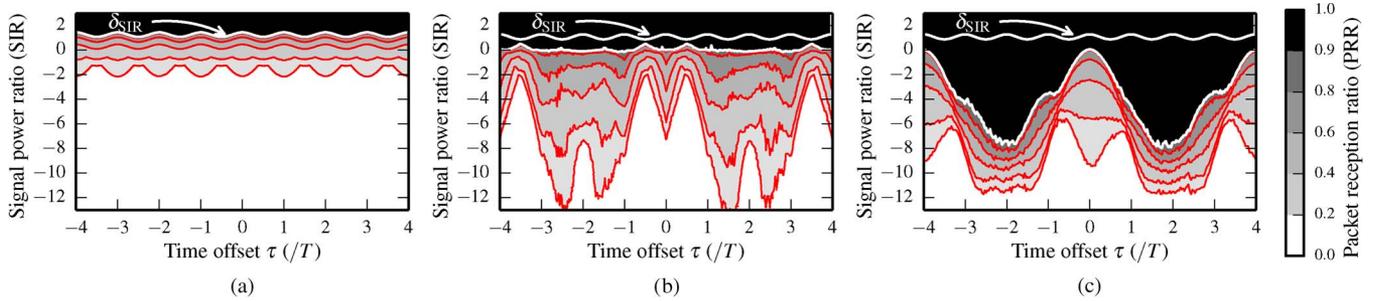


Fig. 4. The capture threshold for two colliding packets with *independent* payload, varying with the signals' power ratio SIR and time offset ($\tau = 0$ indicates that the signals overlap fully). For the uncoded case, the threshold δ_{SIR} is nearly constant across all time offsets and represents the classical capture threshold (thus, for reference, it is drawn in all figures). For HDD, the threshold is nearly constant, but 1 dB lower. Additionally, there is a wide transitional region with non-zero PRRs. Finally, for SDD the threshold is very sensitive to signal timing, we observe a variation of 6–8 dB with periodical time shifts of $2T$. (a) Uncoded transmissions. (b) DSSS with hard decision decoding. (c) DSSS with soft decision decoding.

at the receiver. Our metric of interest is the PRR of the SoI, i.e., we observe the probability to overcome the collision. The results for three classes of receivers are shown in Fig. 4.

Uncoded transmissions: From Fig. 4(a), we observe that the capture threshold is a good model to describe the PRR of interfering, uncoded transmissions. If the SoI is stronger by a threshold δ_{SIR} of 2 dB, all its packets are received.² This behavior persists for all choices of τ , i.e., packet reception is independent from the properties of the interfering signal (we only see a minor periodic effect). Below the threshold, there is a narrow transitional region with non-zero PRR. Under uncoded transmissions, our model is able to recover the classical capture threshold for MSK and is in accordance to experimental results in the literature [6], [8].

Hard decision decoding: When considering HDD (Fig. 4(b)), we note that the threshold abstraction is still valid and the performance improvement of coding is only 1 dB (the coding gain is canceled when the same chipping sequences are used). In the transitional region, there is a wider parameter range that results in non-zero PRRs, e.g., when τ is close to integer values (and thus $\cos \varphi_p \approx 0$), we observe a better PRR for \mathcal{S} . These results show that coding with HDD yields only limited benefits if all senders use identical chipping sequences.

Soft decision decoding: Finally, for SDD we observe a strong dependence between PRR and time offset (Fig. 4(c)). Only for positions without chipping sequence shifts ($\tau = 0$, and because of the way IEEE 802.15.4 sequences are chosen $\tau = 4kT$, $k \in \mathbb{Z}$) the performance is comparable to the HDD case. For different time shifts, we can achieve a 6–8 dB coding gain despite the use of identical chipping sequences; especially for offsets $\tau = 4kT + 2T$, we can achieve a clear coding gain. The reason is that soft bits contain additional information on the detection confidence, which helps to improve the detection performance in the cross-correlation.

This insight suggests that two senders may benefit from coding even when using independent payloads, provided that they time their collisions precisely. This may help to increase the number of opportunities for concurrent transmissions, i.e., interfering nodes can be much closer to a receiver and still achieve the same PRR performance. In other words, a *constant*

capture threshold is too conservative when collision timing can be precisely controlled, because the performance of SDD is very sensitive to time offsets.

2) *Capture Threshold Under Identical Payload:* When considering the collisions of identical packets, we observe very different results (Fig. 5): a good reception performance is possible despite even a negative SIR.³

Uncoded transmissions: For uncoded transmissions, the PRR performance is shown in Fig. 5(a). While in this case the threshold for a PRR of 100% is still equal to the independent payload case, substantially more packets are received in the transitional region with time shifts less than $\pm 0.75 T$. However, PRRs around 30% are usually not sufficient to boost the performance of network protocols. The reason for this limited performance is the carrier phase offset φ_c : with negative SIR, the interfering signal dominates the bit decision at the receiver, and with larger offsets $\varphi_c \in ((\pi/2); (3\pi/2))$, the term $\cos \varphi_c$ changes its sign and flips all subsequent bits. In this sense, the literature conjecture that constructive interference is the reason for the good performance of flooding protocols [17], [18] is only valid if the receiver is synchronized to the strongest signal and if the phase offset φ_c can be neglected. However, because the collisions start during the preamble when using such protocols, successful synchronization cannot be ensured. Therefore, there must be another mechanism that recovers flipped bits.

Hard decision decoding: The reception performance of coded messages provides a hint in this direction (Fig. 5(b)). We observe a corridor of τ values ($\tau = \pm 0.2 T$ or 100 ns in IEEE 802.15.4) that has a PRR of 60–80% in the center (note the larger SIR scale on the y-axis). When two signals with identical payload collide with a small time offset, a reception is still possible even if the interfering signal is far stronger. This suggests that the interfering signal is received instead of the SoI, and that coding helps to overcome bit flips of β_k induced by the carrier phase. The explanation is a property of (8): even if all bits are flipped by $\cos \varphi_c$, the (absolute) correlation is still maximal for the correct chipping sequence. This shows that DSSS used in IEEE 802.15.4 is a key factor to make the collision-aware protocols work.

²For the numerical values of δ_{SIR} shown in the figures, we used a PRR threshold of 90%.

³We note that with increasing time offsets τ the PRR performance approaches the results for independent payloads.

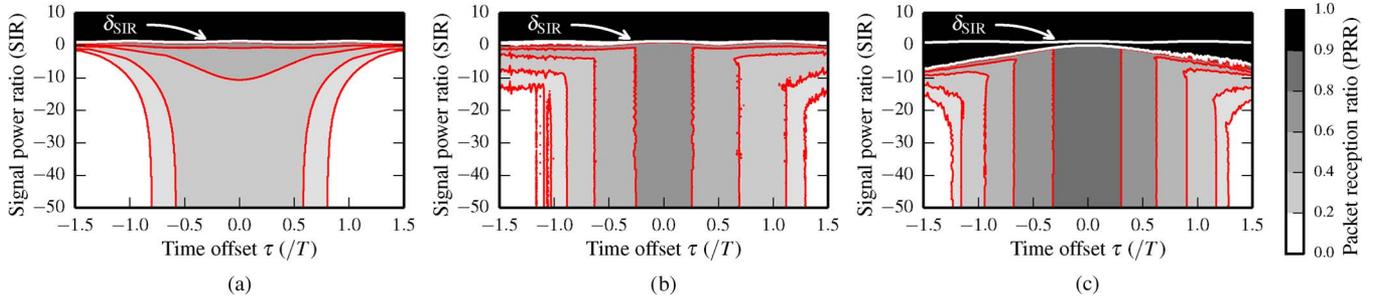


Fig. 5. The capture threshold for colliding packets with *identical* content depending on the power ratio SIR and the time offset τ . In all three figures, we show the threshold δ_{SIR} for identical and uncoded payload as reference. (a) In the uncoded case, the PRR is non-zero in the transitional τ , but packet loss is still likely with PRRs of 20–30%. For coded transmissions, we observe a central area that enables high PRR values (up to 70% in (b) and approximately 90% in (c)). (a) Uncoded transmissions. (b) DSSS with hard decision decoding. (c) DSSS with soft decision decoding.

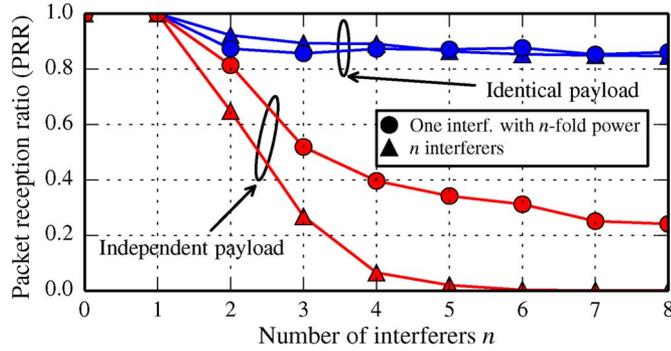


Fig. 6. Reception ratio for SDD under one strong interferer or n weaker interferers, but all with equal received power. For identical payload the difference is small, for independent payload several interferers are more destructive than one.

Soft decision decoding: The experimentally observed performance in the literature is even superior to Fig. 5(b) [11], [15], [17]. Taking SDD into account, this gap is closed (Fig. 5(c)). There is a strong center region for $\tau \leq \pm 0.3 T$, or 150 ns in IEEE 802.15.4, with a PRR of approximately 90%. Now, this matches well with existing experimental results. This means that the reception performance is very good in this center region *independent* of the SIR, i.e., no power control is required and perfect time synchronization is unnecessary for successful reception.

3) *Effect of Several Interferers:* In this subsection, we consider the effect of one strong interferer compared to several interferers with the same power when combined, but evenly distributed across the interferers. We consider the following scenario: all interferers are time-synchronized ($\tau_i = 0$), but each has an i.i.d. uniform random phase offset $\varphi_{c,i}$ (and independent payload bits $\beta_{k,i}$ if different content is assumed). The interference power varies with $(n/2)P_{\text{SoI}}$ for a number of interferers $n \in \{1, \dots, 8\}$, with each interferer having a signal power at the receiver of $(1/2)P_{\text{SoI}}$.

Under the classical capture threshold model both interference types share the same SIR and thus lead to the same PRR at the receiver. However, as we observe in Fig. 6, this is only the case for identical payload, for independent payload n interferers prove to be more destructive despite having the same signal power. While experimental results by Ferrari *et al.* suggested this result [15, Fig. 12] for identical payload, the root cause is now explained by our model: a single interferer is more likely affected by high attenuation ($\cos \varphi_c \approx 0$) than n independent

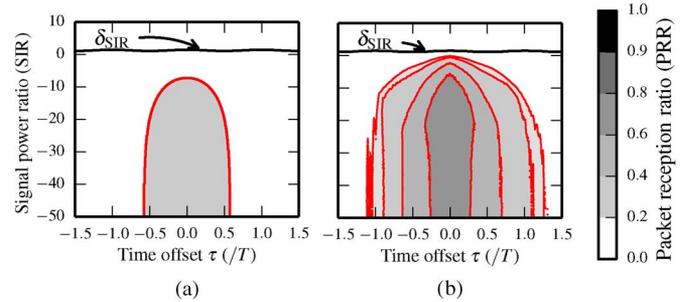


Fig. 7. Reception regions of an interfering signal with *independent* payload for the uncoded and DSSS with hard decision decoding (HDD) case. (a) Uncoded. (b) HDD.

interferers, resulting in a higher likelihood of destructive interference. However, in case of identical payload, even an effective interferer is still received correctly in 90% of the cases. The observation for independent payload reveals another problem of SINR models: relying on the signal power ratio alone discards the crucial effects of each interferer’s offsets.

C. Reception of Interfering Signals With Indep. Payload

Our results explain why and when collision-aware protocols work even without power control: coding enables the reception of interfering signals despite carrier phase and time offsets. In this section, we revisit the case of independent payload but focus our interest now on the reception of the *interfering* signal, i.e., we treat the interfering signal $u(t)$ as the SoI and observe the reception of β_k instead of α_k . Related work by Pöpper *et al.* [21] shows that for uncoded systems the reception of interfering signals is indeterministic; in contrast, we show analytically and experimentally (Section VI) that real systems can receive unsynchronized, interfering packets reliably when using *coded* messages.

Uncoded transmissions: This case is shown in Fig. 7(a). In this case, a reception is only successful if bits are not flipped by either φ_c or φ_p , and we observe a PRR of 20–30% in the center region ($\text{SIR} < -10$ dB and $|\tau| < 0.5 T$) in our evaluation. The reason for the poor reception performance is visible in Fig. 8(a); the acceptable parameter values of τ and φ_c that lead to an error-free packet reception have tight constraints. The interfering signal must hit into a *capture zone* defined by the signal parameters, which permits the signal to have only small time and carrier phase offsets.

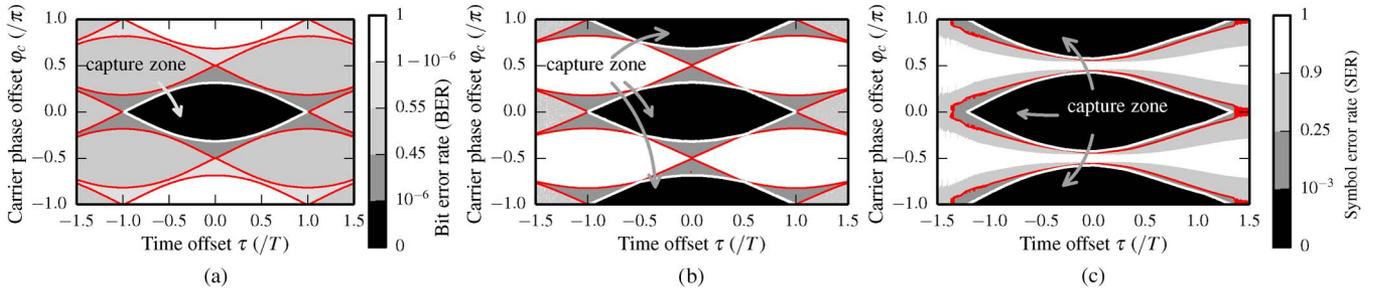


Fig. 8. Relation between error rates and signal parameters, time offset τ and carrier phase offset φ_c (with SIR = -40 dB). An unsynchronized packet is successfully received if the parameter combinations fall inside the dark *capture zones*. (a) For uncoded transmissions, the error rate increases for phase offsets $|\varphi_c| > (\pi/4)$. (b) For coded transmissions and a HDD receiver, the shape of the capture zone is similar to the case in (a), but a second zone around $\varphi_c = \pm\pi$ is present. (c) For coded transmission and SDD, the eye shape is widened, and an increasing number of parameter combinations result in error-free transmissions. (a) Bit error rate for uncoded transmissions. (b) Symbol error rate for DSSS/HDD. (c) Symbol error rate for DSSS/SDD.

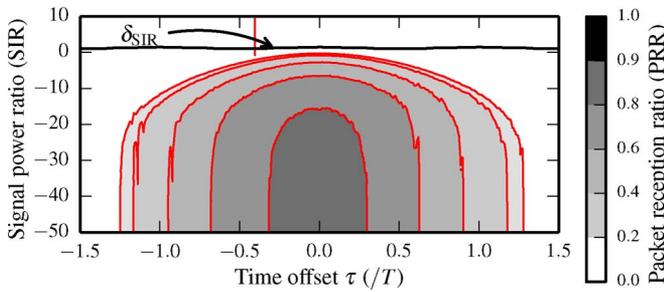


Fig. 9. Reception region of interfering signals $u(t)$ with *independent* payload for the SDD receiver. The threshold for a synchronized signal δ_{SIR} (from Fig. 4(a)) is also shown.

Hard decision decoding: In this setting, the PRR in the central area increases to approx. 60% (Fig. 7(b)). In Fig. 8(b), we see the reason for the increase: while the general shape is the same, we see a second capture zone around $\varphi_c = \pm\pi$. There are two explanations for this. First, we use the same sliced bits from the uncoded case as input for DSSS correlation, which thus possess the same error characteristics. Second, because of the use of absolute correlation values in the correlation (8), the adverse effect of large phase offsets can be repaired. Specifically, this means that even if all bits are flipped, the correlation value is still maximal for the correct chipping sequence. This use of DSSS thus doubles the PRR of an interfering signal.

Soft decision decoding: Finally, in Fig. 9, we see a central area below SIR = -23 dB and a width of 0.25 T that has a PRR for the interfering signal of approx 90%. This means that, if the power difference is large enough, a receiver can ignore a synchronized signal and recover the interfering one despite its offsets. Fig. 8(c) shows this in terms of the capture zone. The eye-shaped regions are much wider compared to the other receiver designs, and especially for the central region with minor deviations of τ , the SER is negligible. Problems in the reception only occur for carrier phase offsets such that $\cos \varphi_c \approx 0$. These results show that interfering signals can indeed be received, which helps in collision-aware protocols or other intentional collisions, e.g., in message manipulation attacks on the physical layer. To validate this new result, we present an experimental study of such receptions with real receiver implementations next.

VI. EXPERIMENTS

In this section, we provide experimental evidence that our model accurately captures the behavior of existing receiver implementations. Since many results in the previous section comply with existing experimental results (see also Section VII), we focus our efforts on the reception of interfering signals because this topic is not well covered experimentally in the literature. We note that we also validated our analytical results with a simulation model based on the numerical integration of time-discrete signals, which confirmed the correctness of our model at the symbol and chip levels. The purpose of this section is to show that our simplifying assumptions, especially for the receiver model, are justified.

A. Experimental Setup

To perform this experiment, the requirements for the interferer differ from the scope of operation of Commercial Off-The-Shelf (COTS) devices. We need to (i) transmit arbitrary symbols on the physical layer, without restrictions like PHY headers, (ii) synchronize to ongoing transmissions with high accuracy, and (iii) schedule transmissions at a fine time granularity. To meet these requirements, we implemented a custom software-defined radio based experimental system.

1) *Interferer Implementation:* To this end, we modified our USRP2-based experimental system RFRReact [40] to recover the timing of the other signal and send arbitrary IEEE 802.15.4 symbols at controlled time offsets. Because of its implementation in the USRP2's FPGA, the system is able to tune the start of transmission with a granularity of 10 ns and send arbitrary waveforms. A detailed description of the system can be found in a technical report [41].

2) *Experimental Methodology:* In our experiments, we consider three parties in the network: a standard-compliant receiver (we monitor the behavior of two implementations to test for hardware dependencies, Atmel AT86RF230 and TI CC2420), a synchronized sender \mathcal{S} (a COTS RZ Raven USB), and the interferer \mathcal{I} described above. The procedure is as follows: \mathcal{S} sends a packet with PHY headers, MAC header, and 8 byte payload. \mathcal{I} time-synchronizes with this signal and schedules the transmission of 8 different bytes at the beginning of the payload of \mathcal{S} . The receiver first synchronizes on \mathcal{S} and receives its

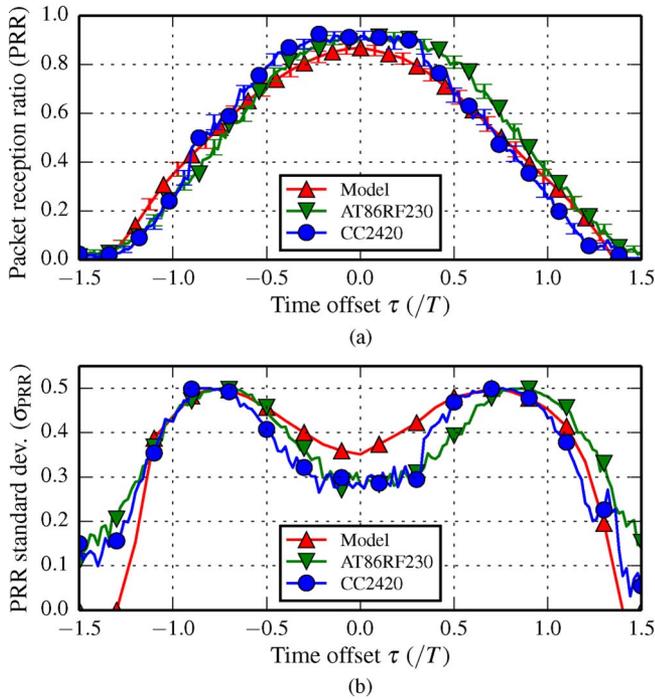


Fig. 10. Experimental results for two receivers in terms of packet reception (PRR) performance and PRR standard deviation compared to our model. Both receivers display a behavior that is well-described by the model. (a) Comparison of packet reception ratios. (b) Comparison of PRR standard deviations.

header, but experiences a collision in the payload bits. We note that the receivers do not attempt to correct bit errors, retransmissions are used for error recovery during normal operation. Damaged packets are simply detected using the checksum at the end and discarded in case of failure. For the experiments we reconfigured the devices so that all packets are recorded, even if the checksums did not match.

We chose values of τ in $(-1.5 T; 1.5 T)$ or ± 750 ns in steps of 10 ns; for each time offset τ , we sent 1000 packets and analyzed the payload detected by the receiver. We derived the value of τ empirically, i.e., we chose the point with maximum PRR in the center as $\tau = 0$. We adjusted the transmit power of \mathcal{I} to result in a SIR of -40 dB to be in the region of interest.

B. Experimental Results

We analyze our measurements using two metrics, packet reception ratio and symbol error rate.

1) *Packet Reception Ratio (PRR)*: Based on the received packet data from the experiments, we derive the PRR as the number of packets with correct payload (of the interferer) divided by the total number of packets. In other words, we measure the empirical success probability for a message manipulation attack. The experimental results for the mean PRR of the two receivers are shown in Fig. 10(a). We observe a good fit with the predictions of our model to both receivers, Atmel AT86RF230 and TI CC2420. In the central region, the receivers show a slightly better ability to receive the interfering signal than predicted by our analytical model. The reason is that our model makes the assumption that no frequency offset is present and that the receiver does not try to resynchronize with

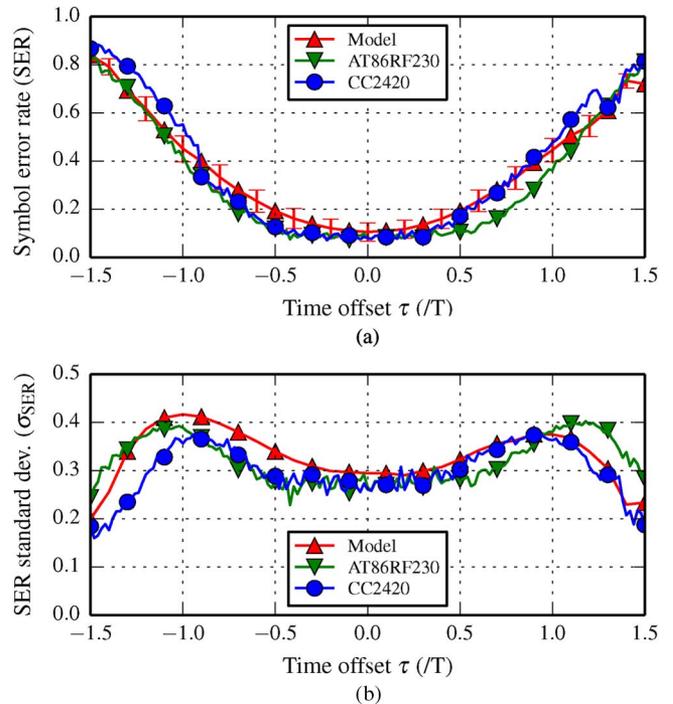


Fig. 11. Comparison of experimentally measured symbol error rates and standard deviations, and the SER values predicted by our model. (a) Comparison of symbol error rates. (b) Comparison of SER standard deviations.

a stronger signal. However, receivers must be able to tolerate frequency offsets of up to 100 kHz [33, Sec. 6.9.4] and thus track and possibly correct the phase during the packet reception process. Yet, as the results show, our assumptions still yield a good approximation of the real receiver behavior.

To further validate our model, we perform an analysis of the standard deviation of the measured PRR values (Fig. 10(b)). In general, the second order statistics follow the non-trivial shape well. On closer inspection, we observe three regions in the graph. For $|\tau| < 0.5$, our model slightly overestimates the standard deviation; the reason is that the PRR performance of the COTS receivers is better than our model, leading to less variance. For $0.5 < |\tau| < 1.1$, the curves are close to each other. Finally, in the zone with $|\tau| > 1.1$, the model slightly underestimates the standard deviation, again because the real receivers perform better than the model predicts. Still, the model provides a good approximation of the behavior of widely used receivers for interfering signals under the assumption of random carrier phase offsets.

2) *Symbol Error Rate (SER)*: We derive the SER by summation of the number of symbol errors across the payload of all received packets for a given time offset τ , and divide this sum by the total number of payload symbols. This metric gives better insights into the causes for packet errors, and provides another validation for the capture zone. In Fig. 11(a), we observe that the fit is good for the symbol error rate as well, with a slightly better SER performance for the COTS receivers as expected. Considering the SER standard deviation (Fig. 11(b)), we observe a similar behavior as in the PRR case, the predictions of the model and the measured results provide a good fit in both, curve shape and absolute values.

VII. GUIDELINES FOR OPTIMAL PARAMETER SETTINGS

Here, we provide a summary of our main findings and highlight key conclusions for the design of protocols that leverage concurrent transmissions. In particular, we summarize how the notion of capture zones enables engineers and protocol designers to choose optimal parameter ranges for signal power ratio, time offset, and carrier phase offset to ensure a successful reception despite collisions.

A. Signal to Interference Ratio SIR

Our model confirms that when the SoI is above the SIR threshold $\delta_{\text{SIR}} \approx 2$ dB, then a successful reception is guaranteed (the capture effect). This is consistent with existing results; for the CC2420 transceiver, Gezer *et al.* [6], Maheshwari *et al.* [7], Dutta *et al.* [11], and Son *et al.* [8] report an experimentally observed threshold of about 3 dB. Considering that their channels were not noise-free and that SINR measurements were collected by the radio transceivers themselves, rather than calibrated measurement equipment such that inaccuracies may arise, this is consistent with our results. If it can be ensured that the stronger signal arrives first and the synchronization process succeeds, the SIR-based capture threshold is a valid model for receiver behavior.

A different matter is the case when the SoI is located in the negative SIR regime, i.e., the interfering signal is stronger than the SoI. This situation occurs if an interferer is closer to the receiver or the synchronization process fails because of a collision during the preamble (which is the case, for example, for the collision-aware flooding protocols). Our model gives better insights in this situation and shows that a reception may still be possible *no matter what SIR*, given that the interfering signal parameters are in the capture zone as defined by the time offset τ and carrier phase offset φ_c . Valid settings for these parameters are discussed below.

B. Time Offset τ

As a guideline derived from the capture zone, the time offset τ should be below $T/2$ for successful concurrent transmissions with identical content, which translates to 250 ns for IEEE 802.15.4. Thus network flooding protocols, for example Glossy, should aim to keep the transmission start time error below this value to ensure a desired PRR above 75%. If $\tau < 200$ ns can be ensured, the achievable PRR is approximately 90%. We note that this ensures worst-case performance (i.e., the SoI is always in the negative SIR regime). The actual performance may be higher in situations with positive SIR or successful synchronization.

C. Carrier Phase Offset φ_c

If the carrier frequency offset at the receiver can be precisely controlled by the senders, there are several options. Interferers can choose $\varphi_c \approx \pm(\pi/2)$ to minimize their effect on the SoI, reducing their influence to signal demodulation. On the other hand, interferers could aim for the capture zone

(e.g., $|\varphi_c| < 0.4\pi$ or $|\varphi_c| > 0.6\pi$ for $\tau = 0$ and SDD) to ensure that their signal is received without errors. There are, however, few approaches in the literature that aim to exploit this. The reason is that the carrier phase at another physical location is hard to predict except in static and free space scenarios because of fading and multipath effects. Pöpper *et al.* [21] show for uncoded QPSK that carrier phase offsets are the major hindrance for a (malicious) interferer to control the bit decisions. In contrast, the results based on our model suggest that such precise phase control is not necessary when DSSS is used, and that intentional message manipulations by deliberate interference are indeed a real threat [42].

D. Number of Concurrent Interferers

Our results in Section V-B3 explain why the number of interferers only has a small impact on reception performance for concurrent transmissions using identical payload. Ferrari *et al.* [15] observed this behavior in their experiments, achieving a stable PRR above 98% for 2–10 concurrent transmissions. Maheshwari *et al.* [7] observed that the SIR threshold is not varying with an increasing number of interferers. On the other hand, Lu and Whitehouse [16] reported a decreasing PRR when the number of interferers is increased. However, the Flash Flooding protocol relies on capture, such that increased time offsets may also influence the results. Some related work claims that a greater number of concurrent transmitters cause problems (Wang *et al.* [17], Doddavenkatappa *et al.* [14]) because “the probability of the maximum time displacement across different transmitters exceeding the required threshold for constructive interference” may increase. Our model shows that these protocol-related issues should be addressed with more precise timing synchronization across the network. For independent payload, we show that 2–3 interferers are sufficient to reduce the PRR significantly. This confirms the effect reported by Gezer *et al.* [6] that the PRR decreases with an increasing number of interferers.

VIII. CONCLUSION

In this article, we presented the first comprehensive analytical model for concurrent transmissions over a wireless channel. As shown in an extensive parameter space exploration, the model recovers insights from experimental results found in the literature and going beyond that, explains the root causes for successful concurrent transmissions exploited in a new generation of sensor network protocols that intentionally generate collisions to increase network throughput or to reduce latency. Our results reveal that power capture alone is not sufficient to explain the performance of such protocols. Rather, coding is an essential factor in the success of these protocols because it crucially widens the capture zone of acceptable signal offsets, increasing the probability of successful reception. Finally, our experimental study of packet reception under collisions shows a good fit and reinforces the validity of our model; as a further contribution, we demonstrated the feasibility of message manipulation attacks over the air experimentally.

REFERENCES

- [1] K. Leentvaar and J. Flint, "The capture effect in FM receivers," *IEEE Trans. Commun.*, vol. COM-24, no. 5, pp. 531–539, May 1976.
- [2] J. Foo and D. Huang, "Multiuser diversity with capture for wireless networks: Protocol and performance analysis," *IEEE J. Sel. Areas Commun.*, vol. 26, no. 8, pp. 1386–1396, Oct. 2008.
- [3] R. Gummadi, D. Wetherall, B. Greenstein, and S. Seshan, "Understanding and mitigating the impact of RF interference on 802.11 networks," in *Proc. ACM SIGCOMM*, Sep. 2007, pp. 385–396.
- [4] A. Kochut, A. Vasani, A. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *Proc. IEEE ICNP*, Oct. 2004, pp. 252–261.
- [5] J. Lee *et al.*, "An experimental study on the capture effect in 802.11a networks," in *Proc. ACM WinTECH*, Sep. 2007, pp. 19–26.
- [6] C. Gezer, C. Buratti, and R. Verdone, "Capture effect in IEEE 802.15.4 networks: Modelling and experimentation," in *Proc. IEEE ISWPC*, May 2010, pp. 204–209.
- [7] R. Maheshwari, S. Jain, and S. R. Das, "A measurement study of interference modeling and scheduling in low-power wireless networks," in *Proc. ACM SenSys*, Nov. 2008, pp. 141–154.
- [8] D. Son, B. Krishnamachari, and J. Heidemann, "Experimental study of concurrent transmission in wireless sensor networks," in *Proc. ACM SenSys*, Nov. 2006, pp. 237–250.
- [9] M. Sha, G. Xing, G. Zhou, S. Liu, and X. Wang, "C-MAC: Model-driven concurrent medium access control for wireless sensor networks," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 1845–1853.
- [10] M. Vutukuru, K. Jamieson, and H. Balakrishnan, "Harnessing exposed terminals in wireless networks," in *Proc. USENIX NSDI*, Apr. 2008, pp. 59–72.
- [11] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis, "Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless," in *Proc. ACM SenSys*, Nov. 2010, pp. 1–14, ACM.
- [12] P. Dutta, R. Musaloiu-E., I. Stoica, and A. Terzis, "Wireless ACK collisions not considered harmful," in *Proc. ACM SIGCOMM HotNets-VII*, Oct. 2008, pp. 4:1–4:6.
- [13] D. Wu, C. Dong, S. Tang, H. Dai, and G. Chen, "Fast and fine-grained counting and identification via constructive interference in WSNs," in *Proc. IEEE/ACM IPSN*, Apr. 2014, pp. 191–202.
- [14] M. Doddavenkatappa, M. C. Chan, and B. Leong, "Splash: Fast data dissemination with constructive interference in wireless sensor networks," in *Proc. USENIX NSDI*, Apr. 2013, pp. 269–282.
- [15] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, "Efficient network flooding and time synchronization with glossy," in *Proc. IPSN*, Apr. 2011, pp. 73–84.
- [16] J. Lu and K. Whitehouse, "Flash Flooding: Exploiting the capture effect for rapid flooding in wireless sensor networks," in *Proc. IEEE INFOCOM*, Apr. 2009, pp. 2491–2499.
- [17] Y. Wang, Y. He, X. Mao, Y. Liu, and X. Li, "Exploiting constructive interference for scalable flooding in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 21, no. 6, pp. 1880–1889, Dec. 2013.
- [18] Y. Wang, Y. Liu, Y. He, X.-Y. Li, and D. Cheng, "Disco: Improving packet delivery via deliberate synchronized constructive interference," *IEEE Trans. Parallel Distrib. Syst.*, 2014, to be published. [Online]. Available: <http://dx.doi.org/10.1109/TPDS.2014.2312198>
- [19] N. Santhapuri, S. Nelakuditi, and R. R. Choudhury, "On spatial reuse and capture in *ad hoc* networks," in *Proc. IEEE WCNC*, Apr. 2008, pp. 1628–1633.
- [20] D. Davis and S. Gronemeyer, "Performance of slotted ALOHA random access with delay capture and randomized time of arrival," *IEEE Trans. Commun.*, vol. COM-28, no. 5, pp. 703–710, May 1980.
- [21] C. Pöpper, N. O. Tippenhauer, B. Danev, and S. Čapkun, "Investigation of signal and message manipulations on the wireless channel," in *Proc. Comput. Security—ESORICS LNCS*, Sep. 2011, vol. 6879, pp. 40–59.
- [22] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," in *Proc. IEEE EmNetS-II*, May 2005, pp. 45–52.
- [23] D. Yuan and M. Hollick, "Let's talk together: Understanding concurrent transmissions in wireless sensor networks," in *Proc. IEEE LCN*, Oct. 2013, pp. 219–227.
- [24] M. Zimmerling, F. Ferrari, L. Mottola, and L. Thiele, "On modeling low-power wireless protocols based on synchronous packet transmissions," in *Proc. IEEE MASCOTS*, Aug. 2013, pp. 546–555.
- [25] K. Jamieson and H. Balakrishnan, "PPR: Partial packet recovery for wireless networks," in *Proc. ACM SIGCOMM*, Sep. 2007, pp. 409–420.
- [26] F. Schmidt, M. Ceriotti, and K. Wehrle, "Bit error distribution and mutation patterns of corrupted packets in low-power wireless networks," in *Proc. ACM WinTECH*, Sep. 2013, pp. 49–56.
- [27] K. Wu, H. Tan, H.-L. Ngan, Y. Liu, and L. M. Ni, "Chip error pattern analysis in IEEE 802.15.4," *IEEE Trans. Mobile Comput.*, vol. 11, no. 4, pp. 543–552, Apr. 2012.
- [28] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 388–404, Mar. 2000.
- [29] P. Cardieri, "Modeling interference in wireless *ad hoc* networks," *IEEE Commun. Surveys Tuts.*, vol. 12, no. 4, pp. 551–572, 2010.
- [30] A. Iyer, C. Rosenberg, and A. Karnik, "What is the right model for wireless channel interference?" *IEEE Trans. Wireless Commun.*, vol. 8, no. 5, pp. 2662–2671, May 2009.
- [31] J. Manweiler *et al.*, "Order matters: Transmission reordering in wireless networks," *IEEE/ACM Trans. Netw.*, vol. 20, no. 2, pp. 353–366, Apr. 2012.
- [32] J. Proakis and M. Salehi, *Digital Communications*, 5th ed. New York, NY, USA: McGraw-Hill, Nov. 2007.
- [33] *Part 15.4: Wireless medium access control and physical layer specifications for low-rate WPANs*, IEEE Std 802.15.4, Sep. 2006.
- [34] S. Pasupathy, "Minimum shift keying: A spectrally efficient modulation," *IEEE Commun. Mag.*, vol. 17, no. 4, pp. 14–22, Jul. 1979.
- [35] T. Schmid, "GNU radio 802.15.4 en- and decoding," UCLA NESL, Los Angeles, CA, USA, Tech. Rep. TR-UCLA-NESL-200609-06, Sep. 2006.
- [36] M. Wilhelm, V. Lenders, and J. B. Schmitt, "An analytical model of packet collisions in IEEE 802.15.4 wireless networks," TU Kaiserslautern, Germany, Tech. Rep. arXiv:1309.4978, Jul. 2013.
- [37] R. K. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, Modeling*. Hoboken, NJ, USA: Wiley, Apr. 1991.
- [38] T. S. Rappaport, *Wireless Communications: Principles and Practice*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, Apr. 1996.
- [39] R. A. Poisel, *Modern Communications Jamming: Principles and Techniques*. Boston, MA, USA: Artech House, Nov. 2003.
- [40] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "WiSec'11 demo: RFReact—A real-time capable and channel-aware jamming platform," in *Proc. SIGMOBILE Mobile Comput. Commun. Rev.*, Nov. 2011, vol. 15, pp. 41–42.
- [41] M. Wilhelm, I. Martinovic, J. B. Schmitt, and V. Lenders, "Air dominance in sensor networks: guarding sensor nodes using selective interference," TU Kaiserslautern, Germany, Tech. Rep. arXiv:1305.4038, May 2013.
- [42] M. Wilhelm, J. B. Schmitt, and V. Lenders, "Practical message manipulation attacks in IEEE 802.15.4 wireless networks," in *Proc. Workshop Proc. MMB*, Mar. 2012, pp. 29–31.



Matthias Wilhelm (S'14) received the M.Sc. degree in computer science and electrical engineering from the University of Kaiserslautern, Germany, in 2009. He is currently pursuing the Ph.D. degree in computer science at the same institution. His research interests include physical layer security for wireless networks, RF jamming and interference, and software-defined radio technology. He was the recipient of the EWSN/CONET M.Sc. Academic Award in 2010 and a winner in the ACM Student Research Competition Grand Finals in 2012.



Vincent Lenders (M'08) received the M.Sc. and Ph.D. degrees in electrical engineering from ETH Zurich, Switzerland, in 2001 and 2006, respectively. He is a Program Manager at armasuisse, where he is leading the Cyberspace and Information Research Program, which focuses on practical aspects of cyber security, sensor networks, and data analytics. Before joining armasuisse, he was a Postdoctoral Research Fellow at Princeton University, Princeton, NJ, USA, in 2007. Since 2012, Dr. Lenders serves as the Industrial Director of the Zurich Information Security and Privacy Center (ZISC) at ETH Zurich, Switzerland.



Jens B. Schmitt received the Ph.D. degree from TU Darmstadt, Germany, in 2000. He is a Professor of computer science at the University of Kaiserslautern, Germany. Since 2003, he has been heading the Distributed Computer Systems Lab (DISCO). His research interests are broadly in performance and security aspects of networked and distributed systems.