Receiver Design with an Adjustable Energy-Signal-Quality Trade-off for IoT Networks

Paul Detterer, *Member, IEEE*, Majid Nabi, *Member, IEEE*,
Hailong Jiao, *Member, IEEE*, and Twan Basten, *Senior Member, IEEE*

Abstract—The energy efficiency of an Internet-of-Things (IoT) receiver can be improved by introducing an adjustable trade-off between signal-quality and energy consumption. In good channel conditions, the receiver can be set to consume less energy per bit, without compromising signal quality in bad channel conditions. We propose a system-level receiver design that enables adequate configuration and combination of signal-quality and energy trade-offs in multiple receiver components. Co-design of all components is essential. We identify the most energy-efficient configurations in our system-level design under different channel conditions. With those configurations, the proposed receiver outperforms a state-of-the-art adjustable receiver with only an adjustable analog front end by several tens of percent in energy per successfully received bit and by 2x in energy-sensitivity configuration range. To show the efficacy of the proposed approach, we integrate a model of the proposed design into the OMNeT++ simulator and show the benefits on an environmental monitoring scenario. In this scenario, we report up to 6x energy savings for the entire transceiver compared to the conventional transceiver design without adjustable receiver.

Index Terms—Low power, IoT, Receiver design, Adjustable sensitivity.

I. INTRODUCTION

Low-rate low-power wireless networks for emerging Internet-of-Things (IoT) applications have rigid energy constraints. This is because the IoT nodes in those applications are typically powered by tiny batteries or through energy scavenging and they are expected to operate reliably for years. In emerging pico-cell multi-hop networks, the receiver becomes a dominant system-level energy consumer. In contrast to reception, required transmission power scales with communication distance, becoming less dominant for short-range communications. Furthermore, the receiver is burdened to provide the synchronization between communicating nodes and therefore is in operation for longer time periods than the transmitter. In this work, we target ultra-low power radios motivated in [1]. It is a challenge to design a radio receiver that meets such energy constraints and still satisfies signal-quality

requirements. Those radios need to operate on extremely low energy budgets for applications with very small node sizes and very long battery-life requirements. To achieve this, further energy consumption reductions are required.

One approach to address the energy-efficiency challenge in IoT receivers is to exploit variations in the channel conditions. As demonstrated by, for instance, [2], the variations in the wireless channel of IoT nodes can be significant. To achieve adjustability to channel variations, different signal-quality vs. energy trade-off configurations can be provided in hardware to match the required signal quality to the transmission and reception efforts. Run-time adaptation has been demonstrated in a variety of receiver components such as the Low Noise Amplifier (LNA) [3], the LNA and Mixer combination [4], the Analog-to-Digital Converter (ADC) [5], and the digital baseband [6]. Recent work demonstrates how to configure an adjustable receiver at network level [7]. Yet, the system-level combination of adjustable components and their impact at the system level is not sufficiently explored. To bridge between network and component levels, we propose a system-level receiver design and an accompanying communication-level model that establishes the relation between receiver sensitivity and consumed energy per bit considering multiple adjustable components. Besides an Adjustable Analog Front-end (AAF) and an Adjustable Digital Baseband (ADB), we include channel coding in our design, in the form of an Adjustable Channel Decoder (ACD), as an effective technique to increase robustness against channel noise. Channel coding enables an error-correction capability at the cost of a decreased payload data rate and additional computational complexity. Through a co-design of AAF, ADB, and ACD, we show that channel coding can be used for an energy efficient reception through trading-off signal quality for energy consumption in the AAF and ADB. This compensates the significant overhead in energy consumption of the channel coding at the receiver side. In short-range wireless communications, that overhead is difficult to compensate in a conventional way through reduced transmitter power, because the transmitter power scales with range and is not a dominant contribution to the total energy consumption. In contrast to conventional use of channel coding, our approach yields significant improvement in energy efficiency of low-power low-range networks and allows integration of channel coding in new applications with strict energy constraints.

The proposed adjustable system is illustrated in Fig. 1. An ADC converts the analog output of the AAF to a digital input for the ADB, in turn connected to the ACD. The adjustable
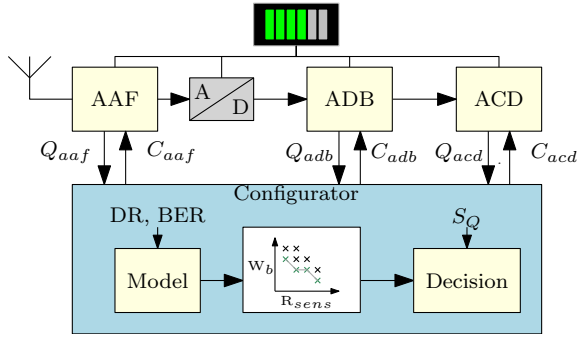
Fig. 1. Receiver design: the signal quality - energy consumption trade-off is adjusted through joint configuration of an Adjustable Analog Front end (AAF, C_{aaf}), an Adjustable Digital Baseband (ADB, C_{adb}), and an Adjustable Channel Decoder (ACD, C_{acd}). The trade-off is estimated through a prediction model that takes communication requirements (Data Rate, DR; Bit Error Rate, BER) as input; configuration decisions are taken based on the observed signal quality S_Q.

elements are configured by a configurator that operates based on a proposed communication model that provides consumed energy per bit (W_b) and receiver sensitivity (R_{sens}) for all system-level configurations ($C = C_{aaf} \times C_{adb} \times C_{acd}$). The model needs the communication requirements such as maximum Bit Error Rate (BER) and the Data Rate (DR) as inputs to estimate the minimum signal-to-noise ratio. With that information, the energy consumption can be reduced when the current observed signal quality (S_Q) is sufficiently higher than the receiver sensitivity R_{sens}. Signal quality can be estimated at the AAF, ADB, and ACD, jointly or separately. Furthermore, the quality can be indirectly estimated at the network level, by tracking the number of received packets, for example. Targeting applications with slow environment dynamics, such as weather changes in environmental monitoring, the signal-quality estimation and adaptation efforts are negligible; we therefore focus on system design and the adjustable components.

Summarizing, the contributions of this paper are as follows:

- We propose a system-level design of an ultra-low-power receiver with adjustable energy vs. signal-quality trade-off. The design integrally covers the analog and digital components of a receiver and outperforms a state-of-the-art adjustable receiver with only an adjustable analog front-end by several tens of percent in energy per successfully received bit and by 2x in energy-sensitivity configuration range.
- We propose an accompanying communication model that makes the sensitivity vs. energy-per-communicated-bit trade-off explicit and a configurator that uses the model to adjust the receiver settings for changing channel conditions.
- We include channel (de)coding to compensate for signal-quality loss in the adjustable receiver front-end, showing that the front-end energy savings can compensate the coding overhead.
- We illustrate the benefits of the approach in an environmental monitoring scenario, showing improvements up to 6x in energy efficiency for nodes operating under good

channel conditions and significant network-level energy saving when compared to a conventional setup without adjustable receivers.

The relevant literature is discussed in the Section II. The proposed communication model is described in Section III. The adjustable receiver design with its configuration options is discussed in Section IV, using the model for motivation and evaluation of choices. The experimental results of using the design as an ultra-low-power IoT receiver are discussed in Section V. This work is summarized in Section VI.

II. BACKGROUND AND RELATED WORK

In the literature, the energy efficiency of a wireless receiver is often addressed at the component level. Conventional system-level receiver design typically focuses on the analog part of a receiver. References [8] and [9] provide good overviews of the state-of-the-art in power-efficient system-level receiver design. For modern digitally intensive receiver designs, the co-design and optimization of the full receiver chain covering both analog and digital components is essential. Moreover, an integral approach to adjustability in the system-level design of a receiver chain provides interesting opportunities for improving energy efficiency. Both the system-level co-design and adjustability aspects are not covered in the literature.

Our design approach is motivated by the insight that for IoT networks, there is a dramatic difference in performance requirements for device operation in the worst-case channel conditions compared to the best case. The traditional worst-case design approach results in poor energy efficiency with varying channel conditions, as noted in [10]. To address that issue, many proposals of adjustable radio elements can be found in literature. An example is the adjustable LNA design described in [3]. The authors highlight that for the majority of time the LNA can be operated with a sensitivity lower than required in the worst case. The authors propose a re-configurable LNA to provide adjustment options in the LNA for better-than-worst-case conditions. In [4], the authors extend the adjustable analog component and include an RF mixer. To address the challenge on the digital side, adjustable digital components are proposed in [10], [11], and [6]. The authors exploit diverse reconfiguration techniques such as under-sampling and approximate computing to enable configurable noise-robustness vs. power-consumption knobs. The mentioned papers thus introduce adjustable component-level knobs that allow to trade signal-recovery efforts for consumed energy. On the network level, [7] shows that such techniques indeed significantly increase energy efficiency.

In this work, we lift the adjustability principle to the receiver system level, bridging the gap between component-level work done so far and the network-level approach. We propose a system-level design of a receiver providing an appropriate combination of adjustable analog and digital elements that enable the reconfiguration of the signal-quality-energy trade-off to meet communication quality requirements at hand, given the channel conditions and incoming signal quality. In contrast to the hypothetic network-level operation modes used in [7],

we provide realistic receiver configurations that are derived from hardware. We evaluate our receiver design in the realistic environment-monitoring network application taken from [12].

The proposed design comes with a communication-level model that relates receiver sensitivity to energy consumption per successfully received bit. This model expands the receiver model described in [13] in several aspects. Compared to [13], we use measurements published on real adjustable components instead of published figures of merits collected from non-adjustable components. Furthermore, we introduce channel coding into the system as a technique to enhance the signal quality. This is a new view on channel coding, which is extensively used to improve the receiver sensitivity and reduce transmission power [14] for transmission systems with non-adjustable receiver components. In such systems, only data rate, number of re-transmissions, and transmission power are adjustable. The integration of a channel coder into a system with adjustable components for re-configurable signal-quality-energy trade-off helps to understand and revise the impact of channel coding on the receiver energy efficiency and allows to use an adjustable receiver in network-level protocols.

III. PROPOSED COMMUNICATION MODEL

To estimate the relation between energy used per successfully received bit and signal quality, we review and modify the conventional communication- and component-level metrics used in receiver and receiver-component design.

The most popular metric used is the receiver sensitivity (R_{sens}). R_{sens} is the minimum signal power at the receiver antenna required to satisfy the maximum Bit Error Rate (BER) constraint. R_{sens} is defined in (1) as described in [15]. It depends on the specific receiver configuration c:

$$R_{sens}(c) = (\frac{S}{N})_{min}(c) \times kT \times BW \times NF(c). \quad (1)$$

$(\frac{S}{N})_{min}(c)$ is the minimum signal-to-noise ratio after the ADC conversion required to meet the maximum Bit Error Rate (BER) constraint. Thus, $(\frac{S}{N})_{min}$ is effectively a function of quality-requirement parameter BER (see Fig. 1), making also R_{sens} a function of BER. Through digital simulations or measurements, the minimum signal-to-noise ratio $(\frac{S}{N})_{min}(c)$ can be estimated for digital components. The term $kT \times BW$ describes the input noise floor at the antenna input. $NF(c)$ is the noise figure, a well-known metric in analog design that describes the signal-to-noise-ratio degradation between input and output of an analog component. For the proposed design, $NF(c)$ and $(\frac{S}{N})_{min}(c)$ depend on the configuration c. Similar to [13], we use the Friis equation [16] to expand the factor $NF(c)$ from (1) to include the contribution of every analog receiver component to the sensitivity.

$$NF(c) = NF_{AAF}(c) + \frac{NF_{ADC} - 1}{G_{AAF}(c)} \quad . \quad (2)$$

$G_{AAF}(c)$ is the overall signal power gain defined in (3).

$$G_{AAF}(c) = \frac{S_{AAF,out}(c)}{S_{AAF,in}} \quad . \quad (3)$$

$S_{AAF,out}(c)$ and $S_{AAF,in}$ are AAF output and input signal powers, respectively.

The noise figure of the ADC, NF_{ADC}, requires special treatment because the converter introduces quantization noise. Assuming white noise and a uniform input-value distribution, following [17], the ADC quantization noise can be described as

$$N_{ADC} = \int_{-\frac{\Delta}{2}}^{\frac{\Delta}{2}} e^2 \times p(e)de = \frac{V_{pp}^2}{12 \times 2^{2n}} \quad , \quad (4)$$

where n is the effective number of quantization bits, Δ is the voltage that corresponds to the smallest ADC output, e is the quantization error, which is always smaller than half of Δ, $p(e)$ is the probability of e to occur, and V_{pp} is the voltage that corresponds to the largest ADC output. Targeting a binary digital number representation, we assume $V_{pp} = 2^n \times \Delta$. As mentioned, the effective number of binary digits (bits) in the digital value representation at the output of the ADC is n. The real number of bits at the output of the ADC is larger than n. However, some of the bits are corrupted by electric circuit noise produced within the ADC. To consider electric circuit noise, the performance of the ADC is matched to performance of an ideal noiseless ADC with n bits [18].

Using (4), NF_{ADC} can be described as

$$NF_{ADC} = \frac{S_{in}/(kT \times BW)}{S_{out}/(N_{ADC} + kT \times BW)}$$
$$= 1 + \frac{V_{pp}^2}{2^{2n} \times 12 \times kT \times BW}. \quad (5)$$

$S_{in}/(kT \times BW)$ and $S_{out}/(N_{ADC} + kT \times BW)$ in (5) are input and output signal-to-noise ratios of the ADC converter, respectively. The ADC provides no signal gain and therefore the ratio of signal powers at input (S_{in}/S_{out}) is one. Combining (1), (2), and (5) yields

$$R_{sens} = (kT \times BW \times NF_{AAF} + \frac{V_{pp}^2}{2^{2n} \times 12 \times G_{AAF}}) \times (\frac{S}{N})_{min}. \quad (6)$$

(6) is used to estimate expected sensitivity for a given configuration c. We omitted c from (6) for the sake of readability.

To estimate energy per bit, we use the quotient of the power sum and data rate.

$$W_b(c) = \frac{P_{AAF}(c) + P_{ADC}(c) + P_{ADB}(c) + P_{ACD}(c)}{R_{uncoded} \times R_{coded}(c)}. \quad (7)$$

The data rate consists of the uncoded data rate $R_{uncoded}$ (in bits/s) and the $R_{coded}(c)$ factor (a number between 0 and 1, depending on c) that describes the data-rate decrease due to additional bits needed for error correction. The data-rate requirement DR in Fig. 1 defines the desired payload rate. The effective payload data rate of a transmitted packet is the product of the uncoded data rate and the coding factor. Therefore, coded and uncoded configurations have a different packet length to achieve the desired payload data rate DR. W_b is thus a function of the DR. $P_{AAF}(c)$, $P_{ADC}(c)$, $P_{ADB}(c)$, and $P_{ACD}(c)$ are power consumption values of the AAF, ADC, ADB, and ACD, respectively.

With (6) and (7), the entire design space of all configurations c of the receiver components can be explored, identifying the most efficient and desirable working points. Each configuration is defined by the combination of the
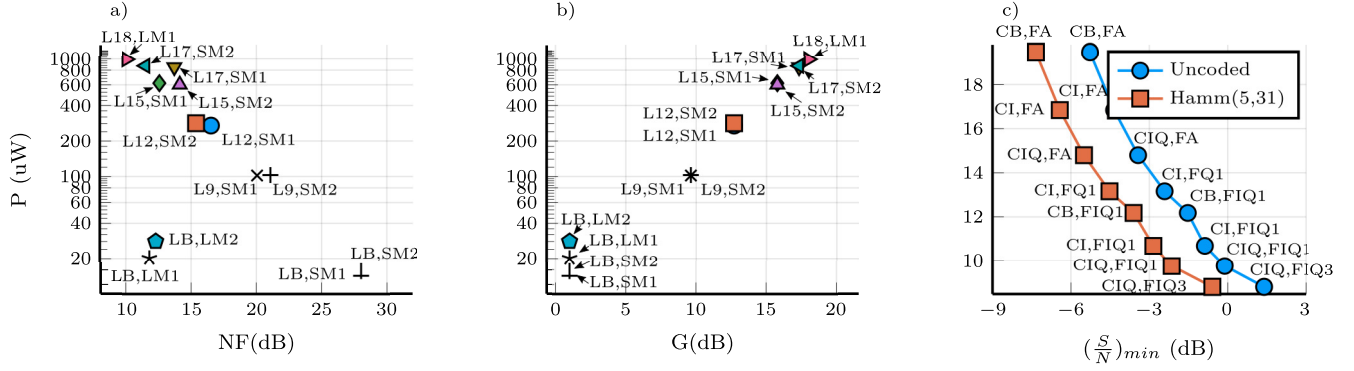
Fig. 2. The component characteristics. The analog estimation of power (P) vs. noise figure (NF), and signal gain (G) are shown in a) and b). The estimation of power consumption and minimum signal-to-noise ratio for coded and uncoded modes are shown in c).

component adjustment parameters (identified in the next section) and the resulting $R_{sens}(c)$ and $W_b(c)$. kT, BW, V_{pp}, n, $(\frac{S}{N})_{min}(c)$, $NF_{AAF}(c)$, $G_{AAF}(c)$, $R_{uncoded}$, $R_{coded}(c)$ $P_{AAF}(c)$, $P_{ADC}(c)$, $P_{ADB}(c)$, and $P_{ACD}(c)$ are the required input to the proposed model and can be estimated through simulation or measurement.

IV. PROPOSED ADJUSTABLE RECEIVER DESIGN

With the model described in Section III, we identify a combination of analog and digital receiver components required for a receiver design with ultra-low-power operation modes below 100 μW. Such a design is beneficial for multi-hop networks used in applications such as environmental monitoring with rigid reception-energy constraints. In this section, we present the overall system design, by describing the essential properties and configurations of individual components in the proposed design and by providing a template configurator component that realizes system-level adjustability. The template configurator component can be further specialized for specific applications, All considered components are implemented to function on a single chip in a commercial 40-nm technology.

A. AAF Design

Considering the power consumption, the Low Noise Amplifier (LNA) is often the dominant part of the receiver analog front end. In LNA-first architectures, the LNA is also an analog component with the most significant impact on sensitivity and constrained by a high bandwidth requirement. The impact of the LNA on sensitivity is twofold. The LNA signal amplification factor G reduces the impact of the internal noise in all following components, as can be seen from Frii's equation [16]. However, a high gain of the LNA also increases its own internal noise and therefore the noise figure of the LNA must be low. The high-gain low-noise-figure LNA circuits are very power hungry. In fact, for the majority of currently used receivers in IoT, such as described in [19], the LNA is the dominant power consumer that defines the limits of IoT-node energy efficiency. For ultra-low-power applications, we see opportunities in a mixer-first architecture where the LNA is placed behind the frequency down-converter (mixer) and baseband filter. Such an architecture exposes a worse noise figure

compared to traditional designs because of significant mixer noise. However, system energy efficiency is better compared to the LNA-first approach because the spectrum of interest at the LNA input is smaller. Therefore, the high bandwidth requirement is reduced. In [4], the authors demonstrate an architecture that allows to bypass the LNA on demand. That architecture allows to use the good signal quality of an LNA-first architecture for poor channel conditions.

The LNA of [4] has a configurable gain providing further noise-power trade-offs. Furthermore, the authors demonstrate ultra-low-power operation in bypass-LNA mode. That mode mimics a mixer-first architecture, and, therefore, is interesting for our work. All configurations of the AAF (C_{aaf}) of [4] are shown in Fig. 2 (a) and (b). The AAF has four bypass-LNA ($G = 1$) modes labeled with prefix LB. The second part of the label refers to a large (LM) or small (SM) mixer switch. The large difference in noise figure between LB,LMi, on one hand, and LB,SMi on the other hand, is because of the better power match between the larger switch and the antenna. The last parameter i indicates the matching strategy described in [4]. The rest of the configuration points are achieved through different LNA configurations that trade internal LNA noise and signal gain for power. That is achieved through multiple switched parallel LNA stages before the mixer switch. The LNA-gain configurations have prefix L followed by the gain.

The authors of [4] do not address the components following LNA and mixer. Considering the noise figure of LNA and mixer only, the LNA-enabled mode of the architecture is inferior to the bypass-LNA mode with a large switch and seems redundant. With our model describing the other elements in the chain, we observe that for $G = 1$ the sensitivity is limited by the following components in the receiver. The reduced signal-power gain due to a bypassed LNA significantly increases the quantization noise inside the ADC converter that limits the gain in noise figure on the analog side. The energy or power savings are also limited by the following components. In bypass-LNA mode, the power consumption of the AAF is so low that it is not dominant anymore on system level. Therefore, to achieve further energy reduction in this mode the co-design with the other components is required. The system-level power consumption can be further reduced through additional trade-off knobs in the digital baseband.

We use the energy efficient implementation of the ADC converter proposed in [20], because this converter has an adequate power vs. bit-width and sampling frequency trade-off. A converter with higher signal quality (through higher bit width or oversampling) would dominate and worsen W_b while a converter with lower signal quality would further limit the performance. As reported in [20], the ADC converter consumes 6 μW with 0.5 V supply voltage and 6.14 MHz sampling frequency. This is three times lower than the power consumption of the used AAF in the lowest power configuration. This is acceptable from the system-level power consumption perspective.

B. ADB Design

To enable further energy reduction, we combine the selected AAF with the ADB proposed in [6]. That ADB uses the efficient demodulation scheme specified by the IEEE 802.15.4 standard [21], and provides multiple configurations to reduce power consumption even further. The configuration points C_{adb} are enabled through variable signal-processing configurations adapting signal bit width and filtering efforts, and applying under-sampling. The ADB has one bypassable digital comparator and two sequential bypassable filtering stages integrated in two parallel computational paths that process even and odd samples. The ADB configurations are shown in Fig. 2 (c). CB, CI, CQ, and CIQ refer to the Comparator Bypassed (CB), Comparator enabled on the In-phase path (CI), Comparator enabled on the Quadrature path (CQ), and Comparator enabled on both paths (CIQ), respectively. FA, FIx, FQy, and FIQz refer to All Filtering stages being enabled (FA), x Filtering stages on the in-phase path I being bypassed (FIx), y Filtering stages on the quadrature path Q being bypassed (FQy), and z filtering stages on both paths combined being bypassed (FIQz), respectively. Some of the possible configurations, including those that under-sample the signal by dropping every other sample, are not shown because they are dominated by others in our design.

C. ACD Design

The analog front end and digital baseband are essential elements of the radio. The channel decoder is of a different nature. It is not essential, but increases the system-level performance while introducing an extra energy overhead. These aspects can be evaluated at system level with our model. As observed from (7), the code rate R_c has a significant impact on the energy efficiency. With power efficient AAF and ADB components, the power consumption of the decoder can easily become an overhead that is not easy to compensate. To be used for energy efficiency, a decoder must not only have high gain in sensitivity but also provide low overhead in power. A decoder suitability metric can be defined as in (8).

$$FOM_{decoder} = \frac{R_{sens}(uncoded) - R_{sens}(coded)}{W_b(coded) - W_b(uncoded)} \quad (8)$$

Though more configurations are possible, we use only two configurations, namely coded vs. uncoded, in our assessment of the use of coding in the ultra-low-power receiver design.
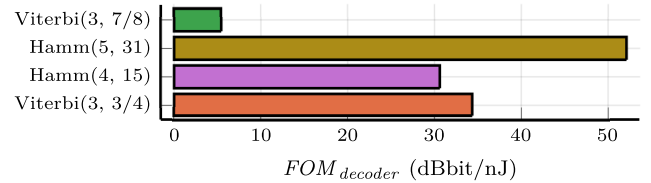

Fig. 3. The decoder efficiency comparison among different decoders.

Due to their high complexity and large power consumption, the state-of-the-art LDPC [22] and Turbo decoders [23] have low $FOM_{decoder}$ and are not suitable for an ultra-low-power receiver. We implemented and evaluated a Viterbi decoder [24] and a Hamming decoder [25] with different parameters. The implementation of the decoders and their evaluation can only be done together with the digital baseband because the error correction performance of a decoder depends on its input distribution. The resulting W_b and R_{sens} for the lowest power LNA-bypass AAF and highest power ADB configuration are shown in Fig. 3. We evaluated two Viterbi decoder configurations with generator polynomial $(7, 5)$ with different puncturing ratios (3/4, 7/8) [24] and two configurations of a Hamming decoder.

Both the Viterbi decoder configurations have constraint length 3 (the first parameter in the figure). Larger constraint lengths are not effective. The Viterbi decoders use a trellis diagram to estimate the likelihood of the received bit stream being one of the valid convolutional code words. The code words are generated through a convolution of the bit stream at the transmitter with the generator polynomial. Decoders with bigger generator polynomials have a too high energy overhead. With a two-element generator polynomial, the code rate of the Viterbi decoder is 1/2, which is too small. Using two puncturing codes from [24], $(6, 5)$ and $(96, 95)$, we increase the code rate to 3/4 and 7/8, respectively.

Compared to Viterbi decoders, Hamming decoders have lower complexity and a significant performance gain. Two variants of the Hamming decoder are compared with two variants of the Viterbi decoder using $FOM_{decoder}$ in Fig. 3, Hamm (5,31) and Hamm (4,15). Hamming codes are systematic codes that introduce smart cross-parity-check bits into the packet providing a Hamming distance of 2 between valid code words [25]. Therefore, Hamming codes can always detect 2-bit errors or correct 1-bit error per code word. Hamming codes have length of $2^n - 1$ with variable parameter n. Within those bits, n bits are parity-check bits and $2^n - 1 - n$ bits are payload bits. The usual referencing to Hamming codes, also used here, is Hamming $(2^n - 1, n)$. The coding rate is $R_{coded} = \frac{2^n - 1 - n}{2^n - 1}$. There is a trade-off between R_{coded} and error-correction capability. We divide the payload of bits in blocks of length $2^n - 1 - n$ and encode them separately. The smaller n, the more blocks and therefore more errors can be corrected for the whole payload. However, the smaller n, the smaller R_{coded}.

The Hamm(5,31) decoder shows the highest decoding efficiency according to (8) and therefore was chosen for our design. To provide robustness against burst errors we added a linear block interleaver. The design of our Hamming (5,31)

decoder anticipates application with IEEE 802.15.4 [21]. This implies a constraint for the maximum packet size of 127 bytes assuming the whole physical packet unit is en/decoded. We fit and interleave 16 Hamming (5,31) code-words into a packet to have 62-byte payload data. The performance of the baseband with the decoder is shown in the orange plot in Fig. 2 (c). The overhead of the Hamming decoder in power consumption is negligible and the receiver only benefits from the increased error robustness. The receiver can operate with the same sensitivity but lower power if the coded mode is used (for example orange CIQ,FA vs. blue CB,FA). This is a property that makes the Hamming decoder very attractive for low-power receiver application. Another attractive property of the Hamming decoder is its high coding rate $R_c = (31-5)/31$. That property helps to keep coding overhead low, also in terms of energy consumption.

D. Configurator

The configurator design is application dependent. The $(W_b(c), R_{sens}(c))$ metric pairs allow tailoring the configuration options to the targeted application context. The difference between receiver sensitivity and transmission power defines the communication range of the node while W_b in combination with the average data throughput defines the battery life. The changes in the channel are tracked through direct or indirect signal-quality (S_Q) measurement. Based on that quality estimation and the $(W_b(c), R_{sens}(c))$ pairs, the reconfiguration is decided in the decision block shown in Fig. 1.

It is important to note the reconfiguration delay and overhead. The ACD is burdened by the highest re-configuration delay and overhead because switching of channel coding needs to be aligned between the receiver and transmitter. In contrast, AAF and ADB re-configurations do not require changes at the transmitter side, provided that the transmission power is sufficiently high. Those can, therefore, be performed on-the-fly (even during a single packet transmission). The delay of AAF reconfiguration depends on the settling time of the analog components while the delay in ADB reconfiguration depends on the total pipeline depth.

Generally, the time variability of the communication channel defines the complexity of the configurator. For environmental monitoring [12], channel variations caused by weather conditions or vegetation are infrequent, resulting in a low reconfiguration frequency (ranging from once every few minutes to once every few hours or even less) and therefore makes reconfiguration overhead negligible. For more dynamic applications, the power consumption of the reconfiguration should be included into the model, as additional power overhead.

Another cost is introduced by the channel sensing. An adequate choice of configuration requires sufficiently precise estimation of current channel conditions. Our approach has unique opportunities because, unlike for transmitter-based techniques such as transmission power or data-rate control, the estimation and adjustment are performed at the same (receiver) node. For transmitter-based techniques, the impact of adjustments can only be measured at the receiver of the signal and therefore measurement is always indirect. Receiver-based adjustment provides direct measurement of the adjustment impact. For long-lasting channel changes such as path loss increase due to rain or vegetation in environment monitoring, the sensing overhead can be reduced to almost zero through utilization of side products of the packet reception, where indicators such as received signal strength are obtained nearly for free.

We elaborate a simple adjustment algorithm to show the efficacy of our approach. For quality measurement S_Q, we use the minimal Received Signal Strength Indicator ($RSSI_{min}$) of a packet and the Received Packet Count (RPC). For environments with low interference, $RSSI_{min}$ is sufficient for estimation of the channel condition. For appropriate signal-quality estimation in environments with high interference, measurements such as Link Quality Indicator (LQI) can be used instead or in combination with $RSSI_{min}$. The set of configuration points (C_{aaf}, C_{adb}, C_{acd}) is chosen according to the (W_b, R_{sens}) Pareto points. This set is referred to as $C_{Pareto}(\mathrm{DR}, BER)$, where we include the quality requirements DR and BER (see Fig. 1) as parameters, to emphasize that these Pareto configurations may change if the communication requirements change (as explained in the previous section). Algorithm 1 shows a program that is executed periodically for a given period, that needs to be aligned with the communication requirement. In our environmental-monitoring experiment, this period is 30 minutes. The algorithm has two internal state variables: the current configuration $c \in C_{pareto}$ and Expected Packet Count (EPC). EPC is required in case there are no packet RSSI measurements. That can happen if the receiver misses all packets because of too low sensitivity. EPC is estimated at the highest sensitivity mode (lines 7, 8) as the count of all packets received since the last update (RPC). Assuming the nodes are sending packets periodically, we expect the nodes to send EPC packets until the next update. To avoid frequent reconfiguration due to fast changes in the channel, a configuration $c \in C_{pareto}$ is chosen with the sensitivity R_{sens} better than $RSSI_{min}$ at least by a margin r. A configuration with a lower sensitivity is set if the minimum packet RSSI measurement since the last update is still larger than the sensitivity of the current configuration $R_{sens}(c)$ plus the margin $2r$ and if no packets are lost ($EPC \le RPC$, lines 10-12). The value $2r$ is chosen to be the trigger threshold for the sensitivity decrease, because then a meaningfully lower sensitivity can be configured that is still larger than the sum of the measured RSSI minimum plus margin r. If packets for some transmitter nodes are likely to be missed, the $RSSI_{min}$ value is compromised and a sufficiently high sensitivity for reception of all packets is unknown, because the packets with lower RSSI were not received. Therefore, if not all packets were received ($EPC > RPN$) the sensitivity configuration is reset to the maximum sensitivity $R_{sens,max}$ (lines 13-15). In the other case, if the measured RSSI minimum is less than the sum of sensitivity and margin, the sensitivity configuration is adjusted to the lowest sensitivity value that is larger than the measured RSSI minimum plus margin margin (lines 16-19). Note that Algorithm 1 does not implement the switching between uncoded and coded configurations. In our experiments, we only use uncoded configurations. Switching between coded and uncoded configurations can be done

Algorithm 1 Algorithm for run-time sensitivity adjustment to conserve energy in long-lasting good channel conditions.

Input:
1: $RSSI_{min}$ ▷ Estimated minimum RSSI (S_Q)
2: RPC ▷ Received Packet Count (S_Q)
3: r ▷ Sensitivity margin
4: $C_{pareto}(\text{DR}, BER)$ ▷ Pareto set of sensitivity modes
States:
5: $c \in C_{pareto}$ ▷ Current receiver configuration
6: EPC ▷ Expected Packet Count
Update: ▷ Executed periodically
7: **if** $R_{sens}(c) == R_{sens,max}$ **then**
8:      $EPC \leftarrow RPC$ ▷ Estimate expected packet count
9: **end if**
10: **if** $(R_{sens}(c) + 2r < RSSI_{min})$ **and** $(EPC \leq RPC)$ **then**
11:      $c \leftarrow \tilde{c} \in C_{pareto}$ with lowest $R_{sens}(\tilde{c})$
12:        s.t. $R_{sens}(\tilde{c}) - r > RSSI_{min}$ ▷ Decrease sensitivity
13: **else if** $EPC > RPC$ **then**
14:      $c \leftarrow \tilde{c} \in C_{pareto}$ ▷ Reset sensitivity
15:        s.t. $R_{sens}(\tilde{c}) = R_{sens,max}$
16: **else if** $R_{sens}(c) + r < RSSI_{min}$ **then**
17:      $c \leftarrow \tilde{c} \in C_{pareto}$ with lowest $R_{sens}(\tilde{c})$
18:        s.t. $R_{sens}(\tilde{c}) - r > RSSI_{min}$ ▷ Increase sensitivity
19: **end if**

Fig. 4. The energy per bit vs. sensitivity trade-off of all the configurations in the proposed design estimated using the proposed model.

similarly to the transmission-power control techniques known from literature. To limit control-packet traffic, in protocols with acknowledged communication, the receiver can initiate the coding switching handshake with extra information in an acknowledgment packet. One bit may be used to indicate whether the next packet is expected in coded or uncoded format. The transmitter sends packets in coded or uncoded format based on the indication from the last received acknowledgment. The sender indicates the used format in the packet header. The latter is needed for the receiver to ensure that the handshake was successfully completed (which may not be the case if an acknowledgment is lost).

Algorithm 1 can be implemented in HW or in SW. To estimate the power consumption overhead, assume it is executed on a commercial micro-controller core such as an Arm Cortex M0 - a very popular choice for low-power IoT nodes. According to the ARM data sheet, at an operating frequency of 16MHz, M0 power consumption is under 100 mW. The algorithm takes around 64 cycles to execute resulting in an energy overhead per update of less than $\frac{100mW \times 64}{16MHz} = 400nJ$. This worst-case estimation can be significantly improved when the implementation is optimized, e.g., by operating at a lower frequency or by an implementation in HW. In addition to the periodic update, we need to keep track of $RSSI_{min}$ and the number of received packets. These updates can be done as part of the packet reception. Minimum and addition operations take a few clock cycles. Assuming 8 clock cycles per received packet, we arrive at a cost of $50nJ$ per received packet. Also this cost can be further optimized. In our experiments in the next section, we conservatively take into account an overhead of $400nJ$ per update of Algorithm 1 and $50nJ$ per received packet.

Note that the presented approach is parameterized by the quality requirements DR and BER. If these requirements change, th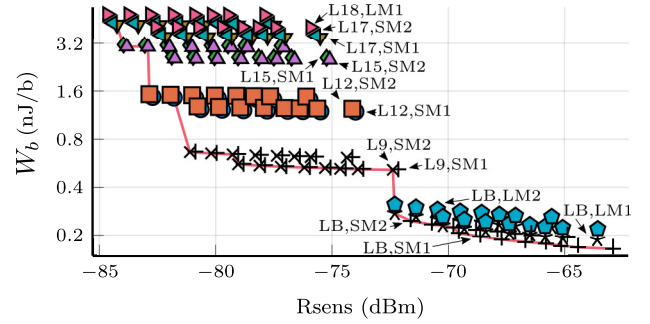en the model elaborated in Section III needs to be recomputed and the Pareto space $C_{pareto}(\text{DR}, BER)$ may change. It may or may not be feasible (in terms of time and energy efficiency) to do such a reconfiguration at run time. This depends on the application and system at hand. In our experiments, we assume fixed requirements.

V. EXPERIMENTAL RESULTS

A. Experimental Setup

To show the advantages of the proposed approach, we apply the proposed model on the proposed design and show the resulting energy vs. signal-quality trade-off. For that purpose, we use the measurements from the AAF reported in [4], the ADC converter described in [20], the ADB reported in [6], and our custom-designed (5,31) Hamming decoder. We use half sine shaped O-QPSK modulation with direct-sequence spread spectrum specified by the IEEE 802.15.4 protocol [21] providing 250 kbit/s data rate in uncoded configuration. The minimum signal-to-noise ratio is estimated as the minimum signal-to-noise ratio needed at the input of the digital baseband for the packet-reception ratio to be above 99% (the IEEE 802.15.4 standard [21]). We use the ADC converter described in [20] with a sampling frequency of 6 MHz, resulting in a power consumption of 6.4 μW and $n = 7$. The values from Fig. 2 are used to compute W_b and R_{sens}. We use Algorithm 1 for switching between sensitivity configurations. As mentioned, we only use uncoded configurations.

To show the efficacy of the proposed approach in a realistic setup, the presented design is integrated in the OMNeT++ network simulator [26] with the results of W_b and R_{sens} computation. The environmental monitoring WSN application reported in [12] is used as a test case for the simulation. Fig. 5 shows the node and routing configuration of the network. In the considered scenario, 52 nodes collect local soil and temperature measurements and send them to a sink node, where the information is processed. The nodes are stationary. The main signal quality metric for this application is Packet Reception Rate (PRR). The target minimal PRR is 99%. Every node sends a packet with 20 bytes of payload once every 15 minutes. Because of the long distances, the packets require several hops to be delivered to the sink node. Some of the routing links are shown in Fig. 5. For example, sensor measurements from node 12 are relayed over nodes 8 and 4 to sink node 0, where the data is aggregated. As a network
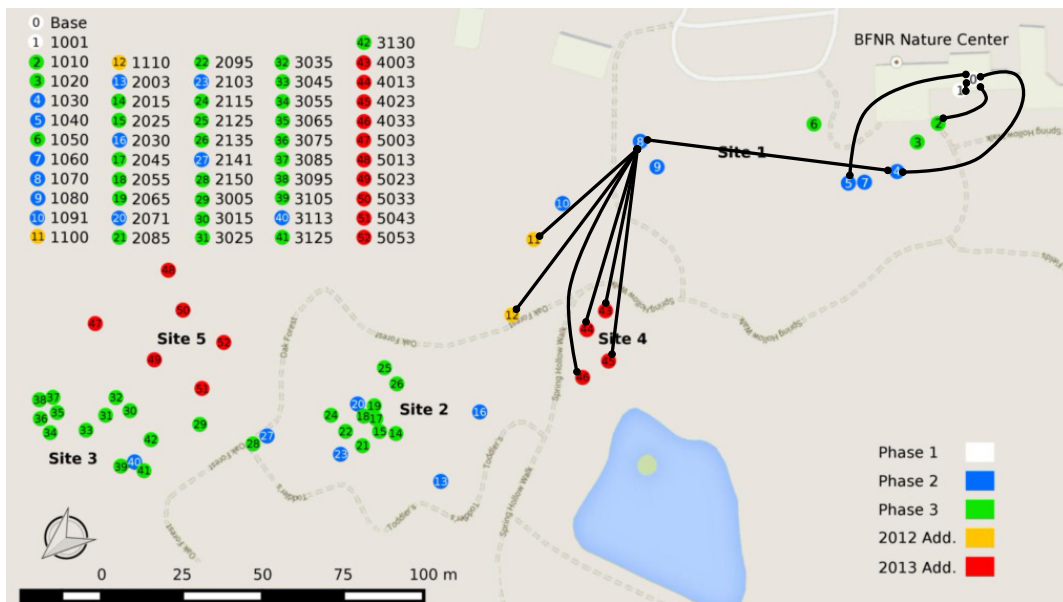
Fig. 5. The environmental monitoring network presented in [12]. With black arrows, some routing links are shown.

protocol, we use non-beacon enabled IEEE 802.15.4 with packet acknowledgments and coordinated sampled listening (duty cycle 1%) [21]. We simulate the given scenario for two days with a 10-hour long 10x increase in path loss that emulates change in channel condition due to rain.

The path-loss model used in simulation is built according to the ITU [27] recommendation and modified by the dynamic path-loss factor $L(t)$ that varies between 0.2 and $10{\times}0.2\ (=2)$ over time as in (9).

$$PathLoss = PathLoss_{ITU}(distance) \times L(t). \quad (9)$$

$PathLoss_{ITU}(distance)$ is estimated using a log normal shadowing path-loss model that accounts for the multi-path fading effect. The background noise power is assumed to be -85dBm. Only intra-protocol interference is considered. Packets with signal to interference-noise ratio below $\left(\frac{S}{N}\right)_{min}$ are considered corrupt and are dropped.

As reference, a receiver with only one sensitivity mode is used (which corresponds to the maximum sensitivity of the adjustable receiver). The reference receiver has no configurator. The configurator overhead of 400 nJ per update and the sensing overhead of 50 nJ per received packet for the configurator implementation as analyzed in the previous section is included in the estimated energy consumption. Updates are made every 30 minutes, i.e., at half the transmission rate.

B. Characterization of the Receiver

All possible energy-sensitivity combinations of the proposed design are shown in Fig. 4. Many configuration points are sub-optimal. The proposed approach enables the identification of the Pareto configurations and drops the inferior ones. The different markers correspond to different AAF configurations from Fig. 2 (a). The AAF is dominating power consumer for all configurations for $G > 1.0$ because of a relatively higher power consumption in comparison to the digital baseband.

This is not the case for $G = 1.0$ configurations as anticipated with the power measurements shown in Fig. 2. Here, the digital components start to be significant and their energy efficiency becomes essential. Fig. 6 shows the zoom-in on the tail of the Pareto front. In uncoded mode and the AAF in the lowest-power configuration (LB,MS1), the ADB provides further energy reduction. This expands the range of possible trade-offs from 10x in sensitivity with the AAF from [4] to 20x with our design.

Another interesting aspect is shown in Fig. 7, showing coded vs. uncoded Pareto-optimal configurations. The coded mode dominates the region of high sensitivity requirements. Compared to an uncoded configuration with similar sensitivity. Up to 25% in energy per bit can be saved by coding. However, from all reconfigurations, it is most difficult to switch between the coded and uncoded modes because this requires reconfiguration of the transmitter side as well. That might not be possible in all network configurations. But the figure shows that within the coded or uncoded mode, also a wide range of set points exists.

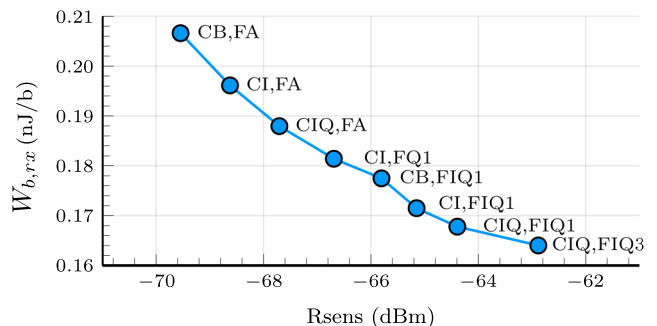With dotted lines in Fig. 7, we show the difference between

Fig. 6. The Pareto-optimal configurations in the lower sensitivity requirement region. The different ADB configurations yield significant energy savings.
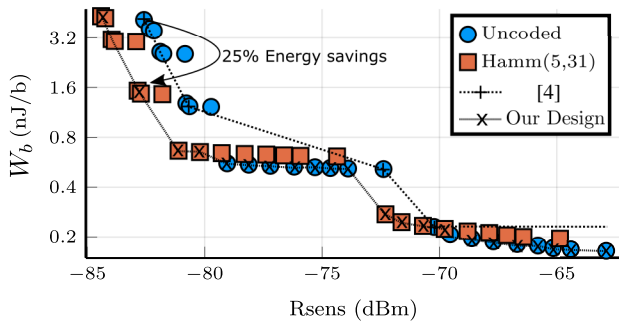
Fig. 7. Coded and uncoded Pareto-optimal configurations.

our design and the state-of-the-art design that includes the AAF from [4] as the (only) adjustable element, combined with the ADB of our design configured for the worst-case channel conditions and no channel coding. Through additional knobs enabled by the proposed design, that includes configurable digital components ADB and ACD, our design outperforms the state-of-the-art adjustable receiver and extends the energy-per-bit/sensitivity configuration range.

C. Network-level Analysis

For the sake of simplicity, we focus on nodes 1, 2, 5, and 8 in the analysis. Node 8 has long communication links as illustrated in Fig. 5. The communication link of node 5 is shorter. The shortest communication links are observed with nodes 1 and 2. Therefore, the energy saving opportunities are the highest for nodes 1 and 2 while node 8 is one of the sensitivity bottlenecks in the network. The energy consumption of the adjustable design and the non-adjustable reference design is shown in Fig. 8. Over time, the nodes consume an amount of energy that is proportional to their momentary power consumption.

Fig. 8 (a) shows the energy consumption of the network nodes without the proposed adjustment. Rain emerges at hour 20 of the simulation, which causes repetitive packet transmission and a tolerable packet loss, resulting in overall packet reception ratio of 99%. The extra energy cost is visible in the slightly increased dissipation rate of, e.g., node 8.

Fig. 8 (b) shows the energy consumption of a network with the proposed receivers. In contrast to the reference, the configurator exploits sensitivity slacks for the nodes 1, 2, and 5 and, therefore, significantly less energy is consumed. From hour 20 onwards, an increase in energy consumption is observed, because the configurator readjusts the sensitivity to handle the higher path loss due to rain.

The sensitivity adjustment is illustrated in Fig. 9. The maximum sensitivity is required at node 8. Therefore, the configurator keeps the sensitivity configuration mostly at the highest level. For nodes 1, 2, and 5, however, the sensitivity can be significantly reduced resulting in energy savings.

After four unsuccessful retransmission attempts, a packet is dropped. The total packet reception rate for the network with receiver adjustment is still around 99%, while the difference with the reference is under 1 %. The re-transmission penalty is more than compensated by the energy savings of lower
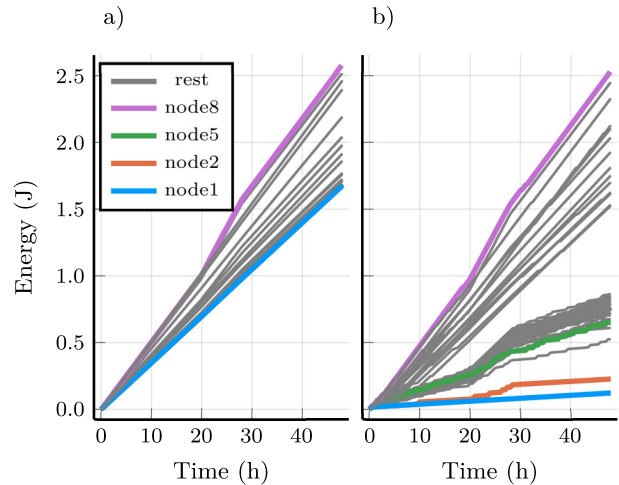

Fig. 8. Energy consumption of the reference design in a) and the adjustable design in b).

sensitivity modes, as observed from Fig. 8. After 48 hours of operation, nodes 1 and 2 consumed over 6x less energy compared to the reference. Node 5 consumed over 2x less energy compared to the reference. The energy consumed by the entire network is shown in Fig. 10. The total energy consumption is reduced by more than 30%.

Note that observed savings are achieved during the rain-free time where slack in sensitivity is large, without any significant loss in packet reception ratio during the rain time. During rain, the sensitivity of many nodes is increased to the maximum and their power consumption (the slope of energy consumption in Fig. 8) and quality-of-service are similar to the reference. The energy savings vary from node to node and depend on whether there is a difference between sensitivity in the worst case and better-than-worst-case channel conditions and time duration of the better-than-worst-case channel conditions. For two rainy days the energy savings for most nodes will be small compared to the reference with fixed maximum sensitivity. Only node 1 with very short distance is still able to operate with low sensitivity and preserve low energy. For two dry days,
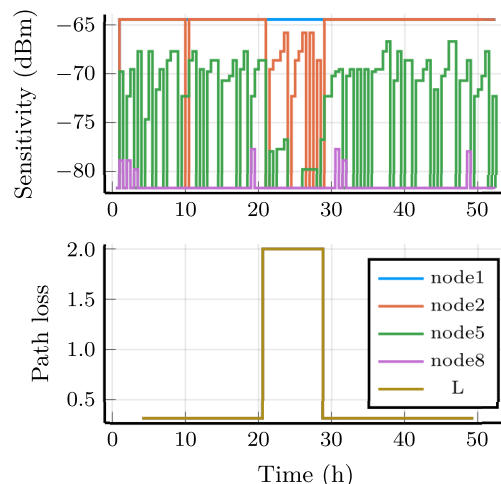

Fig. 9. The sensitivity configuration of the adjustable receivers and simulated path loss (L) over time.
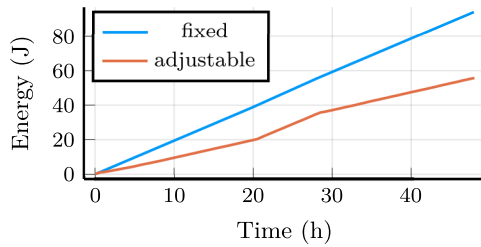
Fig. 10. The total energy consumption of the entire network.

node 5 would consume over 2x less energy compared to the reference. The energy savings for the entire network in dry days are around 50%.

VI. Conclusion

In this work, we present a system-level IoT-receiver design with corresponding model that enables an energy-efficient co-design and configuration of multiple analog and digital components within a low-power wireless receiver. This receiver provides an adjustable signal-quality-power trade-off to achieve a system-level energy-signal-quality trade-off. To limit the channel estimation costs, we reuse information derived from packet reception such as packet RSSI and successfully received packet count. This allows to respond to slow dynamics in the network, like changing weather conditions in environmental monitoring, with negligible overhead. Through the combination of analog and digital adjustable elements and introduction of channel coding, we obtain a more energy efficient design than the design of [4] and extend the range of the energy-quality trade-off from 10x demonstrated in [4] to 20x. Though applied only to an adjustable front end, an adjustable baseband, and an adjustable channel coder, the proposed approach can be extended to an adjustable analog-to-digital converter for more trade-off possibilities. We demonstrate the potential to increase energy efficiency in a realistic environmental monitoring scenario up to 6x for individual nodes and several tens of percent in the overall network compared to the conventional approach that does not exploit the energy-signal-quality trade-off. With the proposed approach, we close the gap in the adjustable-receiver design flow between network-level adjustable receiver algorithms described in [7] and adjustable components proposed in work such as [4] and present a system-level design of an adjustable receiver. The model for signal sensitivity of the radio assumes no side-channel interference and no external interference. This is sufficient for applications with single-channel multiple access. The extension that considers side-channel interference can be achieved by including the non-linearity of analog elements during the frequency conversion, which is interesting future work. The proposed run-time adjustable receiver is complementary to run-time configurable transmitters. The combined network-level adjustment of transmission power and receiver sensitivity in a network is an interesting topic for future work. Also an exploration of run-time receiver-sensitivity adjustment in networks with mobile nodes is an interesting research topic.

References

[1] K. Philips, "Ultra low power short range radios: Covering the last mile of the IoT," in *IEEE European Solid State Circuits Conference*, 2014.

[2] T. Istomin, R. Marfievici, A. L. Murphy, and G. P. Picco, "Trident: In-field connectivity assessment for wireless sensor networks," in *ACM extreme conference on communication and computing*, 2014.

[3] J. Zaini-Desevedavy, F. Hameau, T. Taris, D. Morche, and P. Audebert, "An ultra-low power 28nm FD-SOI low noise amplifier based on channel aware receiver system analysis," *Journal of Low Power Electronics and Applications*, vol. 8, 2018.

[4] G. Chang, S. Maity, B. Chatterjee, and S. Sen, "A MedRadio receiver front-end with wide energy-quality scalability through circuit and architecture-level reconfigurations," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 8, no. 3, pp. 369–378, 2018.

[5] M. Ishizuka, K. Yamada, and H. Ishikuro, "Design of resource sharing reconfigurable $\Delta\Sigma$ SAR-ADC," in *IEEE Asia and South Pacific Design Automation Conference*, 2018.

[6] P. Detterer *et al.*, "Trading sensitivity for power in an IEEE 802.15.4 conformant adequate demodulator," in *IEEE Design, Automation & Test in Europe Conference & Exhibition*, 2020.

[7] P. Detterer, M. Nabi, H. Jiao, and T. Basten, "Receiver-sensitivity control for energy-efficient IoT networks," *IEEE Communications Letters*, 2020.

[8] W. Sheng, A. Emira, and E. Sánchez-Sinencio, "Cmos rf receiver system design: A systematic approach," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 53, no. 5, pp. 1023–1034, 2006.

[9] Z. Fazel, M. Atarodi, and S. Sadughi, "A system-level design method for rf receiver front-end with low power consumption," *Analog Integrated Circuits and Signal Processing*, pp. 1–18, 2021.

[10] S. Dwivedi, B. Amrutur, and N. Bhat, "Power scalable radio receiver design based on signal and interference condition," *Journal of Low Power Electronics and Applications*, vol. 2, no. 4, pp. 242–264, 2012.

[11] V. Lenoir, D. Lattard, and A. Jerraya, "An energy-efficient IEEE 802.15.4 tunable digital baseband targeting self-adaptive WPANs," in *IEEE International Symposium on Circuits and Systems*, 2015.

[12] M. Navarro, T. Davis, G. Villalba, Y. Li, X. Zhong, N. Erratt, X. Liang, and Y. Liang, "Towards long-term multi-hop WSN deployments for environmental monitoring: An experimental network evaluation," *Journal of Sensor and Actuator Networks*, vol. 3, no. 4, pp. 297–330, 2014.

[13] A. Didioui, C. Bernier, D. Morche, and O. Sentieys, "Power reconfigurable receiver model for energy-aware applications," in *IEEE International Midwest Symposium on Circuits and Systems*, 2013.

[14] K. D. Rao, *Channel Coding Techniques for Wireless Communications*. Springer India, 2015.

[15] B. Cook, A. Molnar, and K. Pister, "Low power RF design for sensor networks," in *Radio Frequency integrated Circuits Symposium*, 2005.

[16] H. Friis, "Noise figures of radio receivers," *Proceedings of the IRE*, vol. 32, no. 7, pp. 419–422, 1944.

[17] W. R. Bennett, "Spectra of quantized signals," *Bell System Technical Journal*, vol. 27, no. 3, pp. 446–472, 1948.

[18] S. Haykin, *Digital communication systems*. Wiley, 2014.

[19] Y.-H. Liu *et al.*, "13.2 a 3.7mw-RX 4.4mw-TX fully integrated bluetooth low-energy/IEEE802.15.4/proprietary SoC with an ADPLL-based fast frequency offset compensation in 40nm CMOS," in *IEEE International Solid-State Circuits Conference*, 2015.

[20] K. Yoshioka, A. Shikata, R. Sekimoto, T. Kuroda, and H. Ishikuro, "An 8bit 0.35-0.8V 0.5-30MS/s 2bit/step SAR ADC with wide range threshold configuring comparator," in *IEEE European Solid State Circuits Conference*, 2012.

[21] "IEEE standard for low-rate wireless networks," *IEEE Std 802.15.4-2015 (Revision of IEEE Std 802.15.4-2011)*, pp. 1–709, 2016.

[22] S. Usman and M. M. Mansour, "An optimized VLSI implementation of an IEEE 802.11n/ac/ax LDPC decoder," in *IEEE International Symposium on Circuits and Systems*, 2020.

[23] J. Han, A. Erdogan, and T. Arslan, "High speed max-log-MAP turbo SISO decoder implementation using branch metric normalization," in *IEEE Computer Society Annual Symposium on VLSI*, 2005.

[24] Y. Yasuda, K. Kashiki, and Y. Hirata, "High-rate punctured convolutional codes for soft decision viterbi decoding," *IEEE Transactions on Communications*, vol. 32, no. 3, pp. 315–319, 1984.

[25] R. W. Hamming, "Error detecting and error correcting codes," *The Bell System Technical Journal*, vol. 29, no. 2, pp. 147–160, 1950.

[26] Z. Bojthe, L. Meszaros, and A. Varga. INET Framework. [Online]. Available: https://inet.omnetpp.org

[27] *Propagation data and prediction methods for the planning of short-range outdoor radio communication systems and radio local area networks in the frequency range 300MHz to 100GHz*, ITU Std.
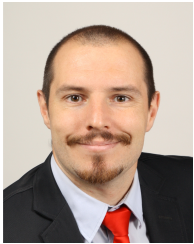
Paul Detterer (S'18-M'22) received the B.Sc in Electrical Engineering, Information Technology and Computer Engineering and M.Sc in Micro- and Nan-otechnologies from RWTH Aachen University in 2016. He is currently guest researcher with Department of Electrical Engineering at TU/e, and research engineer with IC-design group at IMEC Netherlands. His research interests include efficient low-power baseband processing for wireless communication and neuromorphic processors.

Majid Nabi (S'08-M'13) received the B.Sc. and M.Sc. degrees both in computer engineering from Isfahan University of Technology and Tehran University, respectively. He received the Ph.D. degree in electrical and computer engineering from Eindhoven University of Technology (TU/e), Eindhoven, the Netherlands in 2013. He is currently an assistant professor with the Department of Electrical Engineering at TU/e, and Isfahan University of Technology. His research interests include efficient and reliable networked embedded systems, low-power wireless sensor networks, and internet-of-things. He is a member of IEEE.

Hailong Jiao (S'09-M'13) received the Bachelor degree from Shandong University, Shandong, China, in 2004, the Master degree from the Institute of Micro-electronics, Chinese Academy of Sciences, Beijing, China, in 2008, and the Ph.D. degree from the Hong Kong University of Science and Technology, Hong Kong, China, in 2012.

He served as a Tenure-Track Assistant Professor in the Electronic Systems group with the Department of Electrical Engineering, Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands from September 2013 to January 2017. He was with IMEC, Leuven, Belgium as a part-time Visiting Researcher from February 2015 to January 2017. He has been an Associate Professor in the School of Electronic and Computer Engineering, Peking University Shenzhen Graduate School since January 2017. He is also Visiting Assistant Professor with the Electronic Systems group of TU/e. Dr. Jiao's research area is low-power and variations-resilient VLSI circuit and system design.

Twan Basten (M'98-SM'06) received the M.Sc. and Ph.D. degrees in computing science from Eindhoven University of Technology (TU/e), Eindhoven, the Netherlands. He is currently a Professor with the Department of Electrical Engineering, TU/e. He is also a Senior Research Fellow with ESI, TNO, Eindhoven. His current research interests include the design of embedded and cyber-physical systems, dependable computing, and computational models.