

Received 3 October 2023, accepted 14 December 2023, date of publication 22 December 2023, date of current version 11 January 2024.

Digital Object Identifier 10.1109/ACCESS.2023.3346316

RESEARCH ARTICLE

Implementation Framework for Robust Detection, Synchronization, and Decoding of M-ASPM Communications

ALEXEI V. NIKITIN¹, (Member, IEEE), AND RUSLAN L. DAVIDCHACK²

¹Nonlinear LLC, Wamego, KS 66457, USA

²School of Computing and Mathematical Sciences, University of Leicester, LE1 7RH Leicester, U.K.

Corresponding author: Alexei V. Nikitin (avn@nonlinearcorp.com)

This work was supported in part by Pizzi Inc., Denton, TX, USA.

ABSTRACT M-ary Aggregate Spread Pulse Modulation (M-ASPM) is a recently introduced physical layer (PHY) modulation technique that is well suited for use in low-power wide-area networks (LPWANs). Notably, M-ASPM combines high energy-per-bit efficiency, robustness, resistance to interference, and a number of other favorable technical characteristics, with the spread-spectrum ability to maintain the capacity of an uplink-focused network while extending its range. However, while the essential tools for detection and synchronization of pulsed spread-spectrum waveforms in general, and the M-ASPM signals in particular, have been previously provided, a practical framework for combining the detection, synchronization, and decoding of an M-ASPM packet has not yet been suggested. In this paper, we outline such a framework, and describe a prototype algorithm for its implementation. This implementation can be subsequently adapted, under given technical constraints, to specific practical complications such as, for example, significant delay spreads, external technogenic interference, or co-channel and inter-channel collisions. In addition to low latency and computational complexity, the main requirement for this prototype algorithm is that the signal quality remains effectively invariant, for a given path loss, and for a wide range of the data rates, payload sizes, lengths of pulse shaping filters (PSFs), and pulse duty cycles, for a relatively large carrier frequency offset (CFO) between the transmitter (TX) and the receiver (RX). Further, this needs to be achieved without any feedback communications between the TX and the RX, and performed entirely in the RX software.

INDEX TERMS Aggregate spread pulse modulation (ASPM), LoRa, low-power wide-area network (LPWAN), M-ary ASPM (M-ASPM), median tracking filter (MTF), modulo power averaging (MPA), nonlinear signal processing, physical layer (PHY), quantile tracking filter (QTF), spread spectrum.

I. INTRODUCTION AND MOTIVATION

The M-ary Aggregate Spread Pulse Modulation (M-ASPM) is a recently introduced physical layer (PHY) modulation scheme with a number of appealing technical features that have not yet been fully explored and/or quantified. Its first description was given in [1], where the main goal was to examine the spectral and energy efficiencies of coherent and noncoherent M-ASPM variants, and evaluate the bit error probability for M-ASPM links in an additive white Gaussian noise (AWGN) channel. Then in [2] the primary focus was on quantifying M-ASPM networks' scalability, that is, on the spread-spectrum properties of M-ASPM.

The associate editor coordinating the review of this manuscript and approving it for publication was Zhipeng Cai¹.

In particular, we showed that when M-ASPM is used as a spread-spectrum technique (that is, when it operates at the spectral efficiencies below the maximum for a given M , where $\log_2 M$ is the number of bits encoded in a single pulse), its processing gain is proportional to the M-ASPM's average interpulse interval (IpI). As a result, this interval directly affects such link properties as its time-on-air (ToA), the signal-to-interference ratio (SIR) margin and, for a given transmit power, the link's range. In particular, both the ToA and the SIR margin are effectively proportional to the IpI, while the range, if it is a monotonically increasing function of the path loss, is a monotonically increasing function of the IpI. This enables us to maintain the capacity of an uplink-focused M-ASPM network while extending its range.

However, when all M-ASPM nodes transmit with the same average power, implementation of such capacity-preserving range extension may become impractical in complicated propagation environments with greatly varying path losses. Favorably, the efficiency of M-ASPM with constant-envelope pulses can be maintained effectively the same as the efficiency of transmitting a continuous constant-envelope waveform. Then the transmit power of different nodes can be adjusted, without sacrificing the transmission efficiency, to compensate for differences in the path attenuation. This enables us to significantly simplify planning and management of the network. In addition, such a variable-power approach generally increases the network capacity and the average energy efficiency of the nodes, as compared with the arrangement of the nodes with a constant transmit power. Consequently, in [3] we outline a practical approach to implementing such an energy-efficient M-ASPM power control by changing the duty cycle of the pulse shaping filters (PSFs). This can be used for scaling low-power wide-area networks (LPWANs) with realistic desired and/or actual areal distributions of the uplink nodes under diverse propagation conditions.

Nevertheless, even though [1], [2], [3] highlight the main features of M-ASPM communications that make them particularly suitable for LPWANs, they do not provide a holistic roadmap to their practical low-cost deployment. Most importantly, they rely on the assumption that an adequate carrier synchronization is performed beforehand. Since such synchronization at low signal-to-noise ratios (SNRs) is far from trivial, this is an impactful omission that needs to be addressed in practical implementations. In addition, the mathematical description of M-ASPM in [1], [2], and [3] may create an impression that detection of arrival of asynchronous M-ASPM packets in the RX requires continuous filtering, at full sampling rate, with high-order (e.g., with length in hundreds or thousands of samples) finite impulse response (FIR) filters. While for uplink-focused networks this computational burden would be carried by the gateway, it still creates an undesirably costly overhead. Therefore, the goal of this paper is to rectify these omissions by outlining a practical implementation framework for combined detection, synchronization, and decoding of M-ASPM packets.

In our presentation, we simultaneously pursue two main goals. First, we provide a prototype algorithm that can be used “as is” to implement a working M-ASPM system in, e.g., a software-defined radio (SDR). In the description of this algorithm, we do not explicitly include treatment of many important practical complications such as, for example, significant delay spreads, external technogenic interference, or co-PSF and inter-PSF collisions. Instead, as the second goal, in the explanations of the algorithm steps we provide sufficient reasoning and details to facilitate the subsequent adaptations of the algorithm, under given technical constraints, to specific practical scenarios.

For the prototype algorithm, our main target is for the signal quality to remain effectively invariant, for a given path loss, for (i) a relatively large mismatch in the frequency of the local oscillators (LOs) in the transmitter (TX) and the receiver (RX), and (ii) TX and RX motions at relatively high speeds. This should hold for a wide range of the data rates, payload sizes, the lengths of pulse shaping filters (PSFs), and the pulse duty cycles. Further, we want to achieve this invariance without any feedback communications between the TX and RX, any hardware or software changes in the TX, and any hardware adjustments in the RX (e.g., in the LO frequency or sampling time offsets).

As a practical guideline, when choosing the specific numerical values for illustration of the relevant components of the algorithm, we target wireless communications with the nominal carrier frequency $f_c = 915$ MHz, bandwidth $B = 500$ kHz, and ± 30 parts per million (ppm) LO mismatch combined with Doppler shifts due to the relative speeds up to 200 mph.

In Section II, we describe the noncoherent single-sideband M-ASPM with constant-envelope pulses, which is used in the subsequent sections of the paper.

In Section III, we introduce the basic algorithm for synchronization, decoding, and measuring M-ASPM signal quality in an AWGN channel when both the carrier frequency offset (CFO) and the sampling time offset (STO) are negligible, and provide simulated examples of its performance.

In Section IV we show that, when the CFO is still negligible, the impact of a constant STO on the performance of the basic algorithm is insignificant, and illustrate this with simulations.

In Section V, we examine the impact of a CFO on the received pulses, which can be quite significant for long PSFs. Initially, we assume that the cumulative STO due to the associated sampling frequency offset (SFO) is negligible. Then, for the CFO with known sign and magnitude, this impact can be effectively averted by a simple modification of the matched filter in the RX, without introducing any changes to the basic algorithm of Section III.

In Section VI, we evaluate the impact of the cumulative STO due to the mismatch in the frequencies of the LOs, and describe a simple algorithm for its mitigation, which prevents deterioration of the signal quality due to SFO in long M-ASPM payloads.

In Section VII, we demonstrate how the CFO can be accurately measured during the detection of the arrival of an M-ASPM packet. We evaluate the sensitivity of such detection combined with the CFO measurements, and discuss the tradeoffs in matching this sensitivity to the signal quality of the payload. In this section, we also provide the description of the quantile tracking filters (QTFs) used for robust and efficient detection of M-ASPM packets, and assess the properties of the modulo power averaging (MPA) filtering in

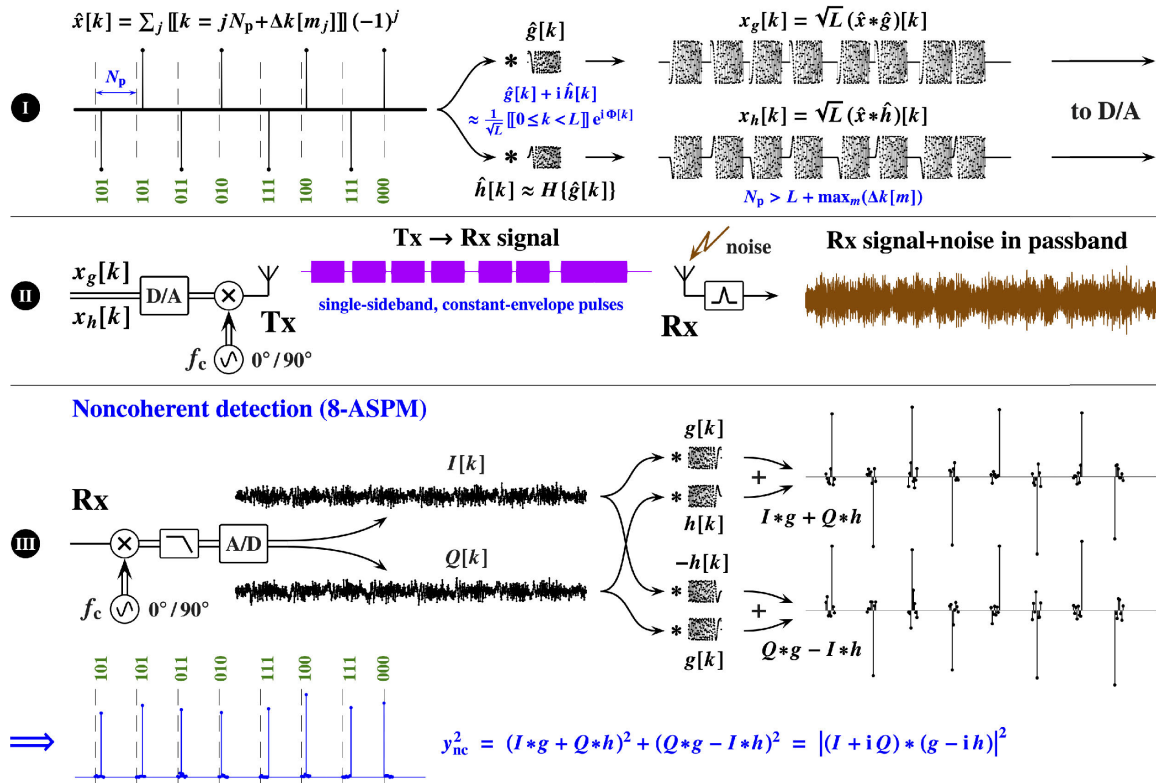


FIGURE 1. Illustration of single-sideband M-ary ASPM link with constant-envelope pulses and noncoherent detection.

its impact on the sensitivity, latency, and the energy efficiency of the detection.

We provide additional comments on the presented approach in Section VIII, along with an example of simulated performance of a full prototype algorithm. In this section, we also address filtering and sampling issues in the context of the computational complexity of the algorithm, and discuss using multiple PSF channels within this framework. We then conclude the paper in Section IX.

Further, all acronyms used in the paper are listed in Appendix A, and the mathematical notations are discussed in Appendix B. Finally, Appendix C illustrates numerical implementation of QTFs.

In several places in the paper, the M-ASPM performance is compared with that of LoRa (short for “Long Range”), a popular modulation technique for LPWANs [4], [5]. For a given number of bits per waveform, LoRa has the same energy-per-bit efficiency as M-ASPM. (In LoRa, this number is represented by the “spreading factor” (SF), which can typically take values from 6 to 12.) Thus, when operating under effectively the same physical conditions (e.g., the same physical frequency band, transmit power, antenna gains, and various system attenuations such as insertion, path, and matching losses, etc.), LoRa represents a suitable benchmark for M-ASPM. When such comparison is given, to distinguish between the respective quantities for LoRa and M-ASPM, we mark those for LoRa by overhead tildes.

II. NONCOHERENT SINGLE-SIDEBAND M-ASPM WITH CONSTANT-ENVELOPE PULSES

For convenience of the reader, let us first briefly describe a particular version of a noncoherent M-ASPM link. This link is illustrated in Fig. 1, and used in the rest of the paper. A more detailed and general M-ASPM PHY description can be found in [1], [2], and [6].

We can encode information in the “arrival times” k_j of the pulses in a digital “pulse train” $\hat{x}[k]$, where only relatively small fraction of samples have non-zero values. Such a “designed” pulse train (an example shown on the left of Fig. 1(I)) can be expressed as

$$\hat{x}[k] = \sum_j \llbracket k = k_j \rrbracket (-1)^j, \quad (1)$$

where k is the sample index, k_j is the sample index of the j -th pulse, and $\llbracket \dots \rrbracket$ is the *Iverson bracket* [7] which is equal to 1 if the expression inside is true and 0 if it is false. The alternating signs of the pulses in (1) simply ensure that $\hat{x}[k]$ is a zero-mean signal. This helps to eliminate a direct current (DC) bias in the modulating signal, which is convenient but not strictly necessary.

For the arrival times in (1) one can use, for example,

$$k_j = jN_p + \Delta N + \Delta k[m_j], \quad (2)$$

where N_p is the average interpulse interval (IpI), ΔN is an integer offset, $m_j \leq M$ is a positive integer, and $\Delta k[m]$ is an integer-valued invertible function such that $0 \leq \Delta k[m] < N_p$

and $\Delta k[m] \neq \Delta k[l]$ for $m \neq l$. The average ‘‘pulse rate’’ f_p in such a train is $f_p = F_s/N_p$, where F_s is the sample rate. For $m_j \in \{1, 2, \dots, M\}$ this pulse train encodes $\log_2 M$ bits per pulse, and thus the raw bit rate f_b is $f_b = f_p \log_2 M$. In the example of Fig. 1, $M = 8$ and $\hat{x}[k]$ encodes 3 bits per pulse. The corresponding 3-bit binary numbers are indicated for each pulse. Note that the peak-to-average power ratio (PAPR) of the designed pulse train $\hat{x}[k]$ is rather large, as it is equal to the IPI $N_p \gg 1$, and this train would be unsuitable for modulating a carrier.

However, the high-PAPR train $\hat{x}[k]$ given by (1) can be ‘‘reshaped’’ by linear filtering, creating a lower-PAPR modulating signal. In particular, the impulse response $\hat{\zeta}_i[k]$ of such a ‘‘pulse shaping’’ filter (PSF) can be a nonlinear chirp with the desired autocorrelation function (ACF), e.g.

$$\hat{\zeta}_i[k] = \hat{g}_i[k] + i \hat{h}_i[k] = \frac{1}{\sqrt{L_i}} \mathbb{[}0 \leq k < L_i] \exp(i \Phi_i[k]), \quad (3)$$

where $\Phi_i[k]$ is the phase and L_i is the ‘‘duration’’ (length) of the chirp in samples. (To generate such a waveform, one can use, for example, the approach described in [8].) In (3), the imaginary part of $\hat{\zeta}_i[k]$ is the discrete Hilbert transform of its real part, i.e., $\hat{h}_i[k] = H \{ \hat{g}_i[k] \}$ [9], [10]. For the i -th PSF $\hat{\zeta}_i[k]$, we will denote its matched filter $\hat{\zeta}_i^*[-k] = \hat{g}_i[-k] - i \hat{h}_i[-k]$ by removing the overhead hat symbol, as $\zeta_i[k] = \hat{\zeta}_i^*[-k]$.

Filtering the designed train $\hat{x}[k]$ with the PSF $\hat{\zeta}_i[k]$ creates the digital modulating signal $z_i[k]$ (‘‘reshaped train’’)

$$z_i[k] = \sqrt{L_i} (\hat{x} * \hat{\zeta}_i)[k] = \sqrt{L_i} \sum_j \hat{\zeta}_i[k-k_j] (-1)^j, \quad (4)$$

where $\hat{\zeta}_i[k]$ is given by (3) and the asterisk denotes convolution. Since in Fig. 1 we show only a single PSF channel, in the figure we omit the subscript i , and also denote the real and imaginary parts of $z[k]$ as $x_g[k]$ and $x_h[k]$, respectively.

After digital-to-analog (D/A) conversion, the real and imaginary parts of $z_i(t)$ can be used for quadrature amplitude modulation of a carrier with frequency f_c , providing the transmitted waveform $\text{Re}(z_i(t)) \sin(2\pi f_c t) + \text{Im}(z_i(t)) \cos(2\pi f_c t)$. Since $\hat{h}_i[k]$ is the Hilbert transform of $\hat{g}_i[k]$, this waveform will occupy only a single sideband with the physical bandwidth B equal to the baseband bandwidth of $\hat{\zeta}_i[k]$ [9]. In addition, if we require that the chirps in (4) do not overlap (i.e., $L_i \leq N_p - \max_m(\Delta k[m])$), then

$$|z_i[k]| = \sum_j \mathbb{[}0 \leq k - k_j < L_i], \quad (5)$$

and, as illustrated in Fig. 1(II), the transmitted signal will consist of constant-envelope pulses. Note that the variance of such a reshaped train is equal to L_i/N_p , and thus, for a given IPI N_p , the average power of $z_i[k]$ is proportional to L_i .

For noncoherent (‘‘nc’’) detection (Fig. 1(III)), in the receiver’s (Rx) quadrature demodulator the noisy passband

signal is multiplied by the orthogonal sinusoidal signals from a local oscillator, lowpassed, and converted to the in-phase and quadrature digital signals $I[k]$ and $Q[k]$. We can then use the matched filters $g[k]$ and $h[k]$, as shown in Fig. 1(III), to obtain the high-peakedness pulse train $y_{nc}[k]$ corresponding to the designed pulse train. Note that after synchronization we would need to obtain only $M = 8$ samples per pulse, i.e., we can use $g[k]$ and $h[k]$ as decimation filters. Out of each 8 samples of $y_{nc}^2[k]$, the position of the sample with the largest magnitude will correspond to the position of the respective pulse in the designed train.

A. UNCODED BER PERFORMANCE OF M-ASPM IN AWGN CHANNEL AND SIGNAL QUALITY CONTROL BY IPI

While AWGN is only a ‘‘background’’ noise in most practical LPWAN applications, the performance in an AWGN channel provides a suitable benchmark for the M-ASPM’s overall efficiency assessment and for examining its main scaling properties.

As a reminder (see [1], [2]), for noncoherent M-ASPM the bit error probability P_b in AWGN channel can be expressed as

$$P_b = P_b \left(\frac{\Gamma}{\eta} \right) = \frac{1}{2(M-1)} \sum_{k=2}^M (-1)^k \binom{M}{k} \exp \left(-\frac{k-1}{k} \frac{\Gamma}{\eta} \log_2 M \right), \quad (6)$$

where $\binom{n}{m} = \frac{n!}{(n-m)!m!}$ is the binomial coefficient, Γ is the SNR, and $\eta = f_b/B$ is the spectral efficiency. The SNR can be further expressed as $\Gamma = (E_b/N_0) \times (f_b/B)$, where E_b is the energy per bit and N_0 is the (one-sided) power spectral density (PSD) of the noise. Notably, as illustrated in Fig. 2, the AWGN bit error probability for M-ASPM is the same as for noncoherent LoRa when $M = 2^{\text{SF}}$, where SF is the LoRa spreading factor [5].

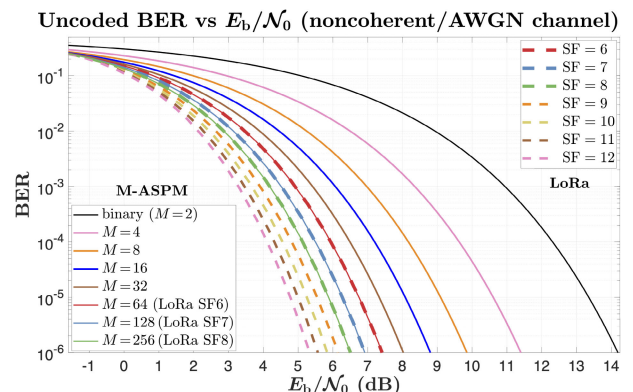


FIGURE 2. Uncoded BER vs E_b/N_0 performances of LoRa (dashed lines) and single-sideband M-ASPM (solid lines) for noncoherent detection in AWGN channel.

If we desire to achieve the same BER performance for the same path loss (i.e., at the same Γ) for LoRa (with a given \tilde{M})

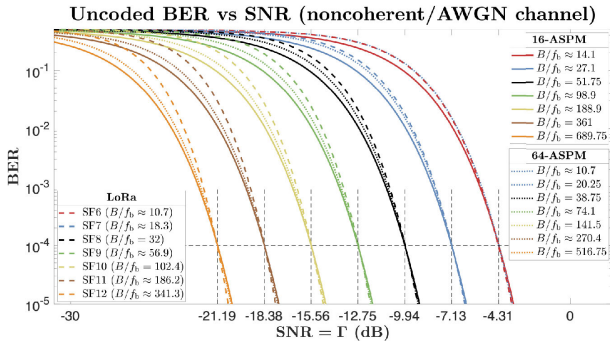


FIGURE 3. Uncoded BER vs SNR performances of LoRa (dashed lines), single-sideband 16-ASPM (solid lines), and single-sideband 64-ASPM (dotted lines) for noncoherent detection in AWGN channel.

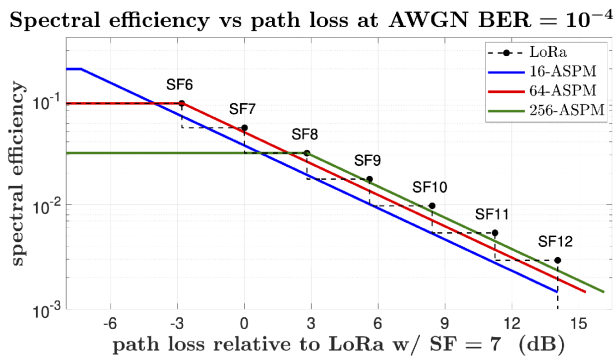


FIGURE 4. Spectral efficiency vs path loss for LoRa and M-ASPM (noncoherent detection).

and M-ASPM (with a given M), the value of η (and thus N_p) can be obtained as a solution of the equalities

$$P_b(\Gamma; M, N_p) = \tilde{P}_b(\Gamma; \tilde{M}) = \text{BER}. \quad (7)$$

An example is given in Fig. 3 for M-ASPM with $M = 16$ and $M = 64$, and $\text{BER} = 10^{-4}$. Then, from the condition $\Gamma/\eta = \text{const}$ it follows that the M-ASPM SNR for a given BER is $\Gamma \propto \eta \propto 1/N_p$. This is in contrast with LoRa, where LoRa’s spectral efficiency is constant for a given spreading factor. While N_p is an integer, it is rather large ($N_p \geq 4M$ for noncoherent M-ASPM) and, for a sufficiently large M (e.g., $M \geq 16$), the M-ASPM’s spectral efficiency can be treated as a continuous quantity. For example, Fig. 4 illustrates M-ASPM’s spectral efficiency vs. path loss at AWGN $\text{BER} = 10^{-4}$, as compared with LoRa. For M-ASPM, for the spectral efficiencies larger than $(\log_2 M)/M$ (i.e., in the spread-spectrum region) the value of η is obtained as the solution of the equality $P_b(\Gamma/\eta; M) = \text{BER}$. For LoRa, the spectral efficiency is the maximum value of $\tilde{\eta}$ satisfying the inequality $\tilde{P}_b(\Gamma/\tilde{\eta}) \leq \text{BER}$.

Without noise, the received pulse train $y_{\text{nc}}[k]$ will be proportional to the convolution of the designed train $\hat{x}[k]$ with the ACF of the PSF. Thus, as discussed in [1] and [2], a good choice for the ACF would be a pulse that combines a small time-bandwidth product (TBP) [11], [12] (e.g., close to that of a Gaussian pulse) with a compact frequency support. An example of such ACF would be a raised-cosine (RC)

pulse [13] with a sufficiently large roll-off factor $0 \leq \beta \leq 1$. Then the sample rate F_s in the digital waveforms can be chosen as $F_s = 2N_s B$, where $1 \leq N_s = 2/(1 + \beta) < 2$ is the oversampling factor. Throughout the rest of this paper, we use $\beta = 1/4$, and thus $N_s = 8/5$.

With this, the value of the spectral efficiency can be obtained as

$$\eta = \frac{f_b}{B} = \frac{2N_s}{N_p} \log_2 M = \frac{16}{5N_p} \log_2 M. \quad (8)$$

Further, the signal quality in the received pulse train $y_{\text{nc}}^2[k]$ can be quantified by the ratio of the peak signal power and the average noise power. This *signal quality indicator* (SQI) λ can be expressed as

$$\lambda = \frac{2\Gamma}{\eta} \log_2 M = \frac{N_p \Gamma}{N_s} = \frac{5}{8} N_p \Gamma, \quad (9)$$

and it is a product of the SNR Γ and the processing gain N_p/N_s .

III. BASIC ALGORITHM FOR SYNCHRONIZATION, DECODING, AND MEASURING SIGNAL QUALITY

In our presentation of the prototype algorithm for detection, synchronization, and decoding of M-ASPM packets, let us begin with an idealized M-ASPM scheme where both the CFO and the STO are negligible, and the noise is the only concern. Further, for the quantitative illustrations we assume an AWGN channel. This basic algorithm is illustrated in Fig. 5. In particular, here we use $M = 16$ and $N_p = 512$, and thus $\eta = 1/40$. The SNR is $\Gamma = -10$ dB, which roughly corresponds to the uncoded AWGN $\text{BER} = 10^{-3}$. The processing gain is $5N_p/8 = 25$ dB, and thus $\lambda = 15$ dB.

A. FIGURE 5(I)

Panel I of Fig. 5 shows the received pulse train $y_{\text{nc}}^2[k]$ obtained by continuous filtering of $I[k] + iQ[k]$ by the matched filter $\zeta[k]$. The information is encoded in non-zero position offsets of the *payload* pulse sequence, highlighted in blue. For the offset $\Delta k[m_j]$ of the j -th pulse, we use $m_j n_{\text{off}}$, with $m_j \in \{1, 2, \dots, M\}$ and $n_{\text{off}} = 8$. For better visibility, we use a relatively short payload sequence with $N_{\text{pl}} = 45$. For synchronization, the payload sequence is preceded by a relatively short *timing* pulse sequence (highlighted in red), where the pulses have zero position offsets. Here we use $N_{\text{sync}} = 9$. In Fig. 5(I), we also show the “noise level” produced as the output $q[k]$ of the median tracking filter (MTF) applied to $y_{\text{nc}}^2[k]$. This output is subsequently used for obtaining both the *robust fence* (threshold) $\alpha[k]$ and the SQI λ .

1) MEDIAN TRACKING FILTER

The MTF output $q[k]$ can be expressed as

$$q[k] = q[k-1] + \mu \operatorname{sgn} \left(y_{\text{nc}}^2[k] - q[k-1] \right), \quad (10)$$

where μ is the *rate parameter*. Further discussion of the *quantile tracking filters* (QTFs), that include the MTF as a

particular case, is provided in Section VII-C and Appendix C of this paper, and more details can be found in [14]. For now, we will note that, for a sufficiently small μ , $q[k]$ will approximate the median of $y_{nc}^2[k]$ obtained in a moving boxcar time window with the width ΔT of order $2 \times \text{IQR}/\mu \gg \langle f \rangle^{-1}$, where $\langle f \rangle$ is the average crossing rate of $y_{nc}^2[k]$ with its median, and IQR is the interquartile range of $y_{nc}^2[k]$.

Our preference of the median over an averaging filter is due to the fact that the former effectively ignores the impulsive M-ASPM signal, and thus $q[k]$ accurately represents the noise level even for high SNRs. In particular, for an AWGN channel, the value of $q[k]$ would correspond to approximately 0.7 times the average noise power. Further, for the same timescale ΔT , $q[k]$ is more *robust* than an output of an averaging filter, as a median filter is insensitive to outliers such as impulsive noise. (Note that even for an AWGN channel the noise is quite impulsive, as it has chi-squared distribution with two degrees of freedom, and its kurtosis is three times that of a Gaussian distribution [15].) A smaller filtering timescale is especially advantageous when the noise is non-stationary, e.g., due to co- and inter-PSF collisions and/or other interference. Also, the computational cost of (10) is even lower than that of the exponential moving average (1st order lowpass IIR filter).

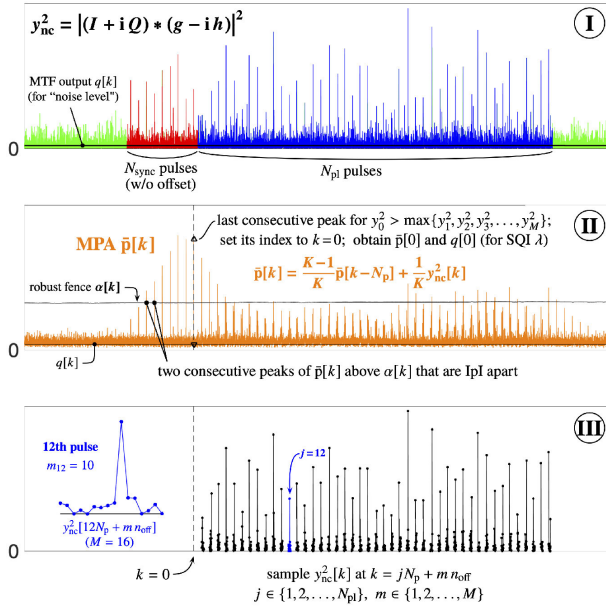


FIGURE 5. Basic algorithm for synchronization, payload sampling, and measuring signal quality when CFO and STO are negligible.

B. FIGURE 5(II)

Panel II shows the output $\bar{p}[k]$ of the *modulo power averaging* (MPA) filter, which can be expressed as

$$\bar{p}[k] = \frac{K-1}{K} \bar{p}[k - N_p] + \frac{1}{K} y_{nc}^2[k], \quad (11)$$

where $K \geq 1$. In particular, $K = 5$ is used in Fig. 5(II). Note that the steady-state output of the MPA filter for a stationary

periodic signal with the period N_p is equal to the input, with the 90% rise time in response to the timing pulse sequence approximately equal to $2.3N_p \times (K - 1)$ (in samples).

We elaborate on the MPA filtering in Section VII. For now, we will note that, without noise, for $K = 5$ and $N_{\text{sync}} = 9$ the peak magnitude of $\bar{p}[k]$ corresponding to the last peak in the timing sequence reaches about 87% of the peak value of $y_{nc}^2[k]$. Also, for $K = 5$ the standard deviation of the noise power is reduced to 1/3 of its average power.

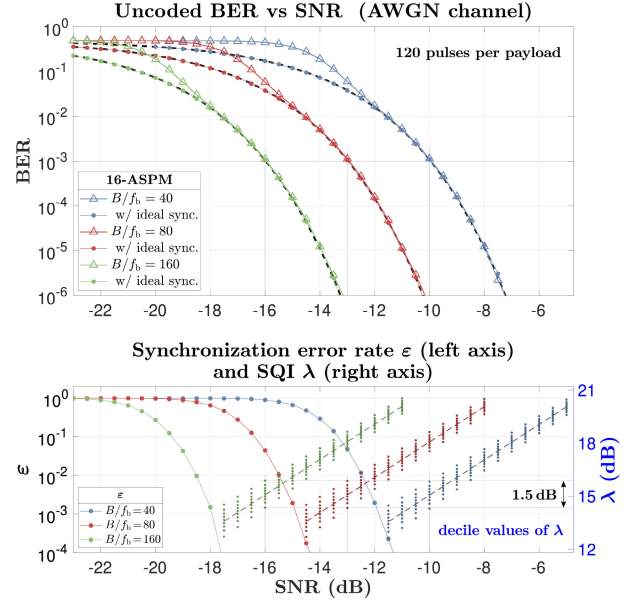


FIGURE 6. Simulated performance of basic algorithm. Black dashed lines in upper panel show theoretical BER with ideal synchronization. In right axis of lower panel, dashed lines are for λ according to (9).

The fence $\alpha[k]$ shown in Fig. 5(II) is simply $\alpha[k] = \gamma q[k]$, where γ is a positive scaling parameter of order ten, and it is “robust” in the sense that it is insensitive to the M-ASPM pulses (as well as to short-duration noise outliers). Specifically, in this illustration $\gamma = 8$. Then, with $K = 5$, $\alpha[k]$ is about 13.8 standard deviations above the average noise power, and the noise peaks that protrude above the fence will be quite rare. At the same time, the “gap” between $\alpha[k]$ and the average noise power is 4.6 times the latter, and, for large processing gains, the magnitude of the signal peaks will significantly exceed this gap even at low SNRs. Then the $\bar{p}[k]$ peaks above $\alpha[k]$ can be reliably detected before the end of the timing sequence.

As a helpful practical tool for detection, timing, and sampling of the pulses in $\bar{p}[k]$, we can use a simple *pulse counting* function $C_\alpha^{\bar{p}}[k]$ that returns unit values for the local extrema of $\bar{p}[k]$ that protrude above the fence $\alpha[k]$, and is zero otherwise. In its essential form, such pulse counting function can be expressed as

$$C_\alpha^{\bar{p}}[k] = \llbracket \bar{p}_k > \alpha_k \rrbracket \llbracket \bar{p}_k > \bar{p}_{k-1} \rrbracket \llbracket \bar{p}_k \geq \bar{p}_{k+1} \rrbracket. \quad (12)$$

For better readability, in (12) we use \bar{p}_k and α_k in place of $\bar{p}[k]$ and $\alpha[k]$.

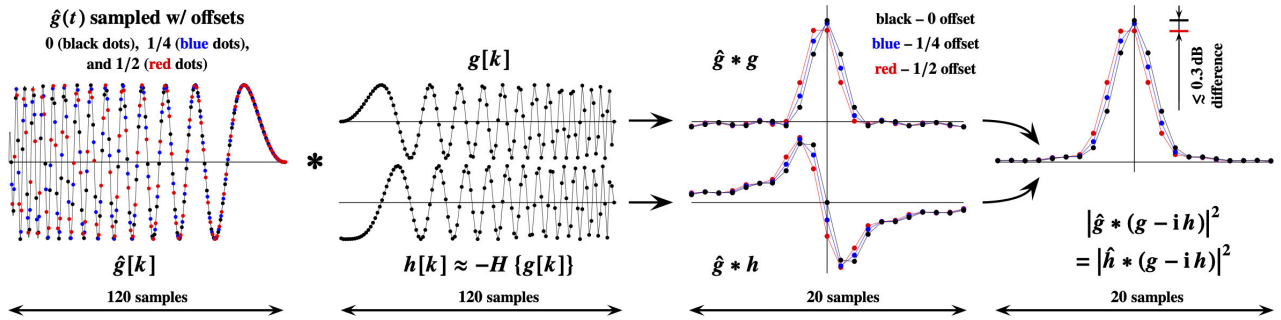


FIGURE 7. Impact of constant fractional sampling time offset on peak pulse magnitude of received pulses is insignificant.

To reduce the number of false-positive detections due to noise outliers (e.g., due to impulsive noise), we further require the detection of two consecutive peaks in $\bar{p}[k]$ that are separated by the IpI N_p . That is, we use the condition $C_\alpha^{\bar{p}}[k_0 - N_p] C_\alpha^{\bar{p}}[k_0] = 1$, where k_0 is the sample index of the second peak.

Next, it is desirable to determine both the sample index and the magnitude of $\bar{p}[k]$ corresponding to the last peak in the timing sequence. The sample index of the last timing pulse provides us with the knowledge of the beginning of the payload, and the magnitude gives us the most accurate measurement of the SQI λ (since this peak contains the largest signal contribution from all peaks in the timing sequence). Therefore, once we obtained k_0 , we evaluate $\llbracket y_0^2 > \max\{y_1^2, y_2^2, \dots, y_M^2\} \rrbracket$, where $y_m^2 = y_{nc}^2[k_0 + nN_p + mn_{off}]$, $m \in \{0, 1, 2, \dots, M\}$, and n is a positive integer. We start with $n = 1$, and increment it until the value of the bracket becomes zero. With this n , $k_0 + (n-1)N_p$ is the sample index of the last timing pulse. For convenience, we reset this index to zero (at the vertical dashed line in Fig. 5(II)).

Then, with $K = 5$ and $N_{sync} = 9$, for AWGN channel the SQI λ can be obtained as

$$\lambda \approx 1.6 \times \frac{\bar{p}[0]}{q[0]} - 2. \quad (13)$$

While, by itself, the knowledge of λ is not necessary for the implementation of any components of the algorithm described in this paper, it is useful for planning and management of M-ASPM networks. For example, it may be wasteful to attempt decoding a long payload if the SQI value obtained during the synchronization is too small. Also, if feedback communications between the TX and RX are available, the SQI knowledge allows us to adjust the TX, RX, and the algorithm parameters in order to optimize the overall performance.

C. FIGURE 5(III)

Once we detected the last pulse in the timing sequence (at $k = 0$), we begin extracting the information encoded in the payload. For this, we sample $y_{nc}^2[k]$ at $k = jN_p + mn_{off}$, where $j \in \{1, 2, \dots, N_p\}$ and $m \in \{1, 2, \dots, M\}$. Then, for each j -th pulse in the payload sequence, the value of m_j can be

determined from the condition

$$y_{m_j}^2 = \max\{y_1^2, y_2^2, \dots, y_M^2\}, \quad (14)$$

where $y_m^2 = y_{nc}^2[jN_p + mn_{off}]$. This is illustrated in Fig. 5(III) for $j = 12$.

D. SIMULATED PERFORMANCE

In Fig. 6, we provide an example of the simulated performance of the basic algorithm described above. It demonstrates that the impact of the synchronization errors on the overall BER efficiency becomes negligible for uncoded BER $\gtrsim 10^{-2}$.

To account for false-negative synchronization errors (“missed packets”), the BER for the respective “missed” payloads (120 pulses each) in the simulations is assumed to be at its maximum (BER = 1/2). This way, the impacts of the false-positive and false-negative synchronization errors on the average BER are effectively equivalent, and there is no need to discriminate between the two. Hence, in the left axis of the lower panel, we plot only the total synchronization error rate ε .

In the right axis of the lower panel, the dashed lines represent the SQI calculated as $\lambda = 5N_p\Gamma/8$ (see (9)), and the dots plot, for each SNR value in the simulations, the decile (10)-quantile values [16] of λ obtained according to (13) (for 40,000 packets at each SNR). One can see that such measured values of λ provide meaningful assessment of the signal quality even for weak signals. For example, even for relatively large BER = 10^{-3} , 80% of the measured values of λ lie within about $\pm 17\%$ of the value calculated according to (9).

IV. INSENSITIVITY TO CONSTANT SAMPLING TIME OFFSET

Let us now assess the impact of a constant STO on performance of the basic algorithm when the CFO is still negligible.

Clearly, the synchronization procedure described in the previous section makes any constant integer STO irrelevant. However, fractional STOs still need to be considered, as they alter the peak magnitude of the received signal pulses. Favorably, as illustrated in Fig. 7, this magnitude is not significantly affected by fractional STOs, and its variation

is confined to about $\pm 3\%$ range. As further illustrated by the simulations shown in Fig. 8, this translates into the respectively small uncertainties in the RX sensitivity and/or measured SQI.

Therefore, the impact of a constant STO on the performance of the basic algorithm is insignificant. Nevertheless, in Section VI we will revisit the impact of a STO for *variable* (cumulative) STOs due to the sampling frequency offset (SFO).

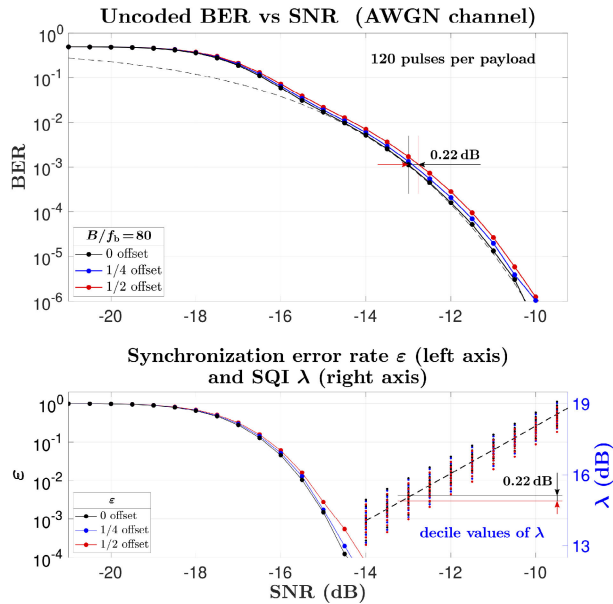


FIGURE 8. Simulated performance of basic algorithm for constant fractional STOs (0, 1/4, and 1/2).

V. IMPACT OF CARRIER FREQUENCY OFFSET AND ITS COMPENSATION

When considering the impact of the carrier frequency offset (CFO) Δf_c on a single received pulse, we can usually ignore the SFO. Indeed, even for the ± 30 ppm SFO, the cumulative change in the STO will remain below 0.3 for PSFs of lengths smaller than 10, 000. Therefore, in practice, the SFO's impact on a single pulse will be negligible.

If we use the value $3/(4\pi^{1/2}\Delta f_c)$ for the coherence time [17], then the condition for the CFO Δf_c impact on a received pulse to remain “small” can be expressed as

$$|\Delta f_c| \lesssim 0.423 \frac{F_s}{L} \approx 1.354 \frac{B}{L}, \quad (15)$$

where L is the length of the PSF. For the CFOs with larger magnitudes, the deterioration of the received pulse becomes significant. This is illustrated in the left panel of Fig. 9.

Favorably, if the value of Δf_c is known, this deterioration can be effectively reversed by using, instead of $\zeta[k]$, the adjusted matched filter $\zeta_\Delta[k]$ in the RX, where $\zeta_\Delta[k]$ is expressed as

$$\zeta_\Delta[k] = \zeta[k] \exp\left(-i2\pi \frac{\Delta f_c}{F_s} k\right). \quad (16)$$

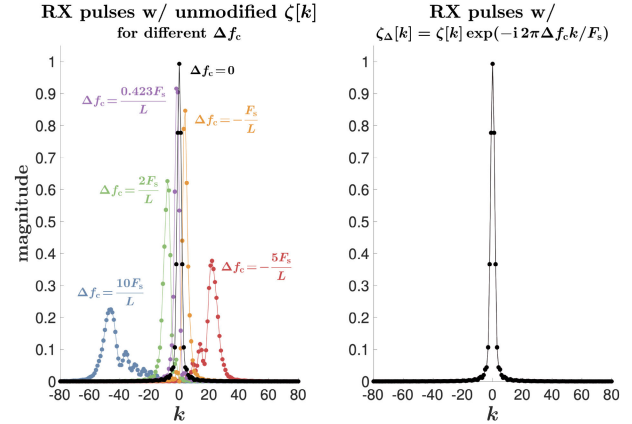


FIGURE 9. Illustration of impact of different CFOs on received waveform for unmodified (left) and adjusted (right) matched filters. For adjusted matched filters all received pulses are identical.




FIGURE 10. Example of matched filter modification to compensate for CFO Δf_c .

The effectiveness of such CFO compensation is illustrated in the right panel of Fig. 9, and an example of the impulse responses of $\zeta[k]$ and $\zeta_\Delta[k]$ is given in Fig. 10.

Note that the time shifts shown in the left panel of Fig. 9 are for an upchirp $\zeta[k]$ (downchirp PSF $\zeta[k]$). These shifts will be opposite in sign if $\zeta[k]$ is a downchirp (upchirp $\zeta[k]$). This difference in time shifts for “flip” PSFs can be used for development of various enhancements and extensions of the prototype algorithm for packet detection and CFO measurements described in Section VII.

A. LO MISMATCH AND DOPPLER SHIFT COMPENSATION

If the mismatch in the frequencies of the LOs is small (e.g., within ± 3 ppm), then our ± 30 ppm targeted range of the CFO tolerance enables effective Doppler shift compensation for significantly large relative speeds. For example, the magnitude of relative Doppler shift for a satellite at the altitude 200 km above the ground remains below about 26 ppm (7.9 km/s speed), and the rate of change of this shift is within ± 1 ppm/s [18]. Then this Doppler shift, combined with the ± 3 ppm LO mismatch, falls within the ± 30 ppm range.

On the other hand, for the relative speeds up to 200 mph (< 90 m/s) the Doppler shift remains below ± 0.3 ppm, which is less than 1% of our targeted CFO tolerance. Then the CFO

can be assumed to be almost entirely due to the ± 30 ppm LO mismatch.

VI. OFFSET IN SAMPLING FREQUENCY AND COMPENSATION FOR CUMULATIVE STO

If the LO frequencies of the RX and the TX are f_c and $f'_c = f_c + \Delta f_c$, respectively, then the ratio of the sampling frequencies is $F'_s/F_s = 1 + \Delta f_c/f_c$.

For a single pulse, the resulting impact of this SFO on the STO would normally be negligible. For example, even for the ± 30 ppm mismatch of the LOs (i.e., $\Delta f_c/f_c = \pm 3 \times 10^{-5}$), the cumulative change in the STO for a PSF of length $L = 1,000$ would be only $\mp 3\%$, and the STO for the entire pulse duration can be assumed constant.

However, the cumulative STO difference δk_j between two pulses separated by jN_p samples is

$$\delta k_j = j \delta k_1 = -j N_p \frac{\Delta f_c}{f_c}, \quad (17)$$

and it can become quite significant for a sufficiently large number j of the IpIs. For example, for $N_p = 2,000$ and the ± 30 ppm mismatch, $|\delta k_1| = 0.06$ and the magnitude of δk_j will exceed half of the sampling interval when $j > 8$.

For a noise-free signal, the impact of such cumulative STO on the magnitude of the sampled pulses in the payload is illustrated in the upper panels of Figs. 11 and 12, for different values of the STO for the first pulse. Here, and throughout this section, in the RX we use the modified matched filter $\zeta_\Delta[k]$ given by (16), and thus the $y_{nc}^2[k]$ pulses remain narrow.




FIGURE 11. Impact of cumulative STO on magnitude of sampled payload pulses (upper panel), and its compensation to within ± 1 (lower panel), for different initial STO values. TX frequency is smaller than RX frequency ($\Delta f_c < 0$).

A. STO COMPENSATION TO WITHIN ± 1

Let us define the *nearest integer function* $\text{nint}(x) = \lfloor 1/2 + x \rfloor$, where $\lfloor x \rfloor$ is the floor function. Then, by adding $\text{nint}(j \delta k_1)$ to the sampling indices of the j -th pulse, we can keep the impact of the cumulative STO to within ± 0.5 samples from the STO of the first pulse in the payload. In other words, we sample the payload at

$$k = jN_p + n_{\text{off}} m + \text{nint}(j \delta k_1), \quad (18)$$




FIGURE 12. Impact of cumulative STO on magnitude of sampled payload pulses (upper panel), and its compensation to within ± 1 (lower panel), for different initial STO values. TX frequency is larger than RX frequency ($\Delta f_c > 0$).

where δk_1 is given by (17) and $m \in \{1, 2, \dots, M\}$.

With this, the total STO for any pulse is confined to the ± 1 interval, and the respective deterioration in the magnitude of the sampled peaks in the payload remains below about 1 dB. This is illustrated in the lower panels of Figs. 11 and 12.

B. STO COMPENSATION TO WITHIN ± 0.75

Of course, the impact of any particular STO of a pulse can be accurately mitigated by employing a matched filter that is sampled with the respective offset δk_j . This is illustrated in Fig. 13 for the fractional STO 1/2. Such matched filters resampled for fractional STO corrections can then be used, in combination with the integer adjustment in the sampling index described above, to further reduce deterioration in the signal quality due to cumulative STO.

For example, in Fig. 14 we alternate $\zeta[k]$ and $\zeta_+[k]$ for the payload filtering, confining the total STO for any pulse to within ± 0.25 samples from the STO of the first pulse in the payload. Thus the total STO for any pulse is confined to the ± 0.75 interval, and we lose less than about 0.5 dB of the signal quality relative to the first pulse.

VII. PROTOTYPE ALGORITHM FOR PACKET DETECTION COMBINED WITH CFO MEASUREMENTS

Implementation of the CFO and the SFO corrections described in Sections V and VI requires knowledge of the CFO Δf_c , which needs to be obtained ahead of processing the timing pulse sequence. Therefore, we need to develop a procedure for obtaining the CFO value at the beginning of an M-ASPM packet. In this section, we introduce such a procedure, that combines the M-ASPM packet detection with the CFO measurements.

First, let us consider a “regular” (i.e., without pulse-position offsets) designed pulse train given by

$$\hat{x}[k] = \sum_{j=0}^{j_{\text{max}}} \llbracket k = jN_p \rrbracket (-1)^j. \quad (19)$$




FIGURE 13. Correction of fractional STO = 1/2 by resampled matched filter $\zeta_+[k]$.

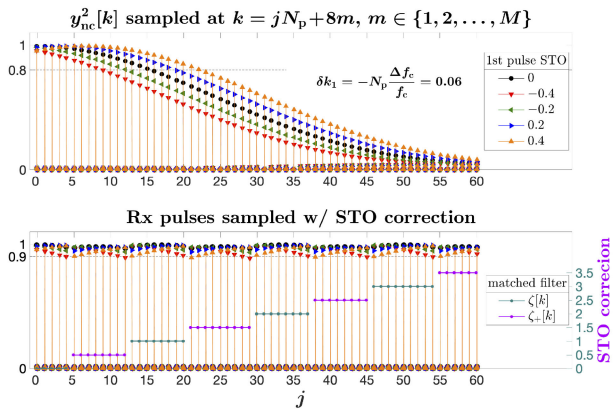


FIGURE 14. Impact of cumulative STO on magnitude of sampled payload pulses (upper panel), and its correction to within ± 0.75 (lower panel), for different initial STO values. TX frequency is smaller than RX frequency ($\Delta f_c < 0$).

For this train, if the cumulative STO is negligible, the following relation will hold for the values of the noise-free received signal sampled at $k_j = jN_p$:

$$y_{nc}[k_j] = y_{nc}[0] (-1)^j \exp\left(-i 2\pi \frac{\Delta f_c}{F_s} k_j\right), \quad (20)$$

where $j \in \{0, 1, \dots, j_{\max}\}$. From (20) it follows that, by measuring $y_{nc}[k]$ at the indices corresponding to the peaks in $y_{nc}^2[k]$, we can determine Δf_c in the $\pm F_s/(2N_p)$ range with $F_s/(2N_p j_{\max})$ accuracy.

To be able to measure a given CFO Δf_c using (20), the value of N_p needs to be sufficiently small, namely

$$N_p \leq \frac{F_s}{2|\Delta f_c|} = \frac{8B}{5f_c} \frac{f_c}{|\Delta f_c|}. \quad (21)$$

For example, $N_p \leq 29$ for the ± 30 ppm CFO, the nominal carrier frequency $f_c = 915$ MHz, and the bandwidth $B = 500$ kHz. Favorably, if the IpI satisfies (21), then for a PSF with the length $L \leq N_p$ the condition $|\Delta f_c| \leq F_s/(2L)$ necessarily holds, and implementing a CFO correction described in Section V is not necessary.

A. LEADING PULSE SEQUENCE

Therefore, a regular pulse train with a sufficiently small IpI, shaped with a sufficiently short PSF, can be used for measuring the CFO. Consequently, we can use such a train as a *leading* pulse sequence, preceding the timing sequence

in an M-ASPM packet, for detection of the packet's arrival and for obtaining the CFO value. Then this CFO value can be used to modify the RX's matched filter for the CFO compensation (see Section V), and to implement the STO correction described in Section VI.

An example of such leading pulse sequence is provided in Fig. 15. Here we use $L' = N'_p = 29 \ll L$, and the leading sequence is a constant-envelope signal with the same amplitude as the pulses in the timing and the payload sequences.

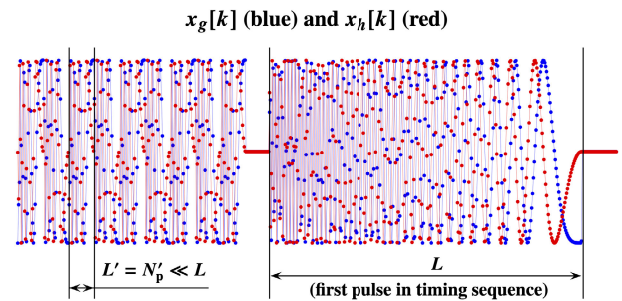


FIGURE 15. In leading pulse sequence we use short PSFs and small IpIs.

B. PACKET DETECTION COMBINED WITH CFO MEASUREMENTS

Let us now demonstrate how a leading pulse sequence can be used for packet detection combined with the CFO measurements. Initially, for visual clarity, we will consider an idealized noise-free case.

First, let us examine the waveforms shown in Fig. 16. Here, the "primed" quantities relate to the output obtained by filtering $I + iQ$ with $\zeta'[k]$, which is the matched filter for the PSF $\hat{\zeta}'[k]$ used to create the leading pulse sequence. For the numerical values, we use $L' = N'_p = 29$ (12.6 dB processing gain) and $N_p = 512$ (25 dB processing gain). Further, the pulse duty cycle in the timing and payload sequences is about 70%, and the CFO is within $\pm F_s/60$ range.

In Fig. 16(I), we show the RX signal $y_{nc}^2[k]$ obtained by filtering $I + iQ$ with $\zeta_\Delta[k]$. As one can see, the peak magnitude of $y_{nc}^2[k]$ in response to the leading pulse sequence is orders of magnitude smaller that the magnitude of the pulses in the timing and the payload sequences. Thus the presence of the leading sequence has no impact on the basic algorithm for synchronization, decoding, and measuring the

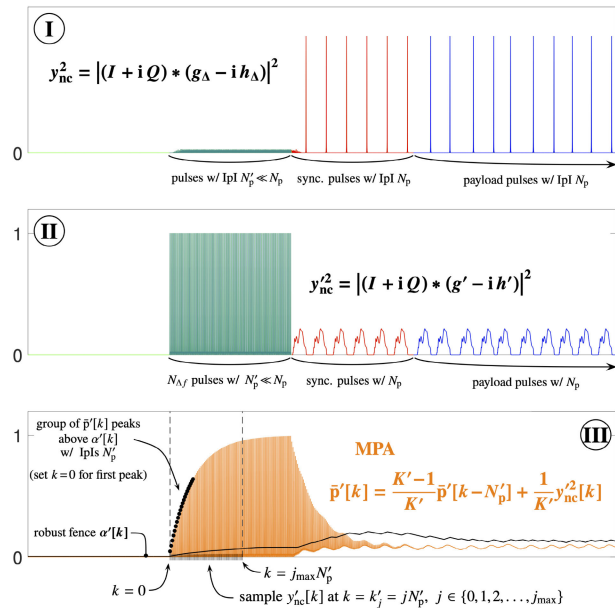


FIGURE 16. Received waveforms without noise for packet detection combined with CFO measurements.

signal quality. In other words, if we start applying $\zeta_{\Delta}[k]$ (to obtain $y_{nc}^2[k]$) before the end of the leading sequence, the performance of the basic algorithm will not be affected.

While the processing gain for N_p' is much smaller than for N_p , $y_{nc}^2[k]$ corresponding to the leading sequence still consists of well-defined, relatively narrow pulses. This is shown in Fig. 16(II). In particular, the 1st quartile of $y_{nc}^2[k]$ for the leading sequence is less than 0.25%, and the IQR is less than 2.5%, of the peak pulse magnitude. Thus, by sampling $y_{nc}^2[k]$ at the indices corresponding to the peaks in $y_{nc}^2[k]$, from (20) we can reliably determine, with $F_s/(58j_{max})$ resolution, Δf_c that falls within the $\pm F_s/58$ range.

In Fig. 16(II), the number of pulses in the leading sequence is rather large, $N_{\Delta f} = 106$. However, the time duration of the leading sequence is still relatively small, as it spans only six IpIs N_p .

The MPA output $\tilde{p}'[k]$ shown in Fig. 16(III) is obtained for $K' = 21$, and the robust fence $\alpha'[k]$ is constructed using the quantile tracking filters described in the next subsection. For reliable detection of the onset of $\tilde{p}'[k]$ in the presence of noise, we require that relatively many peaks of $\tilde{p}'[k]$ (21 peaks in this illustration) with IpIs N_p' protrude above $\alpha'[k]$. (This is in contrast with only two consecutive peaks in the basic algorithm.) We elaborate on this in the next subsection.

Once such group of peaks is detected, we determine the sample index of the first peak. Then, for convenience, we reset this index to zero (at the left vertical dashed line in Fig. 16(III)), and obtain the values of $(-1)^j y_{nc}^2[j N_p']$, where $j \in \{0, 1, \dots, j_{max}\}$. Here, $j_{max} = 64$. These values are plotted in the left panel of Fig. 17 for several different CFOs.

Further, the right-hand panel shows the periodograms of $(-1)^j y_{nc}^2[j N_p']$ [19], where the locations of the peaks correspond to the measured CFOs. Here, the shown physical

frequency range is for the nominal carrier frequency $f_c = 915$ MHz and the bandwidth $B = 500$ kHz.

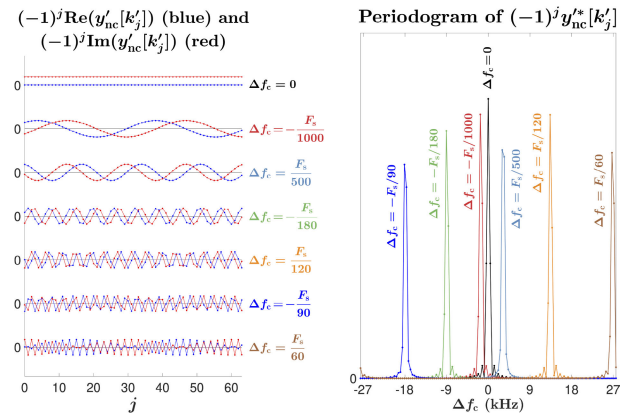


FIGURE 17. Measured CFOs without noise ($k_j' = j N_p'$).

C. MPA FILTERING, AND ROBUST FENCING WITH QUANTILE TRACKING FILTERS

As M-ASPM operates at small spectral efficiencies (e.g., in 10^{-1} to 10^{-3} range), the noise power would be typically more than an order of magnitude larger than the signal power (e.g., $\Gamma \lesssim -10$ dB). Further, for $N_p' \ll N_p$, the processing gain for the leading pulse sequence is much smaller than that for the payload (e.g., only 12.6 dB for $N_p' = 29$). Thus even the peak power of the received pulses can remain smaller than the average noise power.

Favorably, the steady-state output of the MPA filter for a stationary periodic signal with the period N_p' is equal to the input, with the 90% rise time in response to the leading pulse sequence approximately equal to $2.3N_p' \times (K' - 1)$ (in samples). At the same time, as illustrated in Fig. 18, the dispersion of the noise is inversely proportional to $\sqrt{K'}$ for large K' . Thus, for a sufficiently large K' , we can detect the signal's peaks as protrusions above the fence $\alpha'[k]$ that is somewhat above the mean value of the noise. Consequently, for establishing such robust fence $\alpha'[k]$ in the presence of dominant noise, we need the means to continuously monitor both the central tendency and the dispersion of the noise. (Of course, the noise will be zero-mean if the MPA filtering is applied to the difference $y_{nc}^2[k] - y_{nc}^2[k - \Delta N_p']$, where, e.g., $\Delta N_p' = \lfloor N_p'/2 \rfloor$, and we would only need to monitor the noise dispersion. However, such simplification would come at the expense of doubling the noise variance. Nevertheless, this approach should be considered in the subsequent development, in particular, when considering the effects of co- and inter-PSF collisions.)

As a practical matter, quantile tracking filters (QTFs) [14], [20], [21] are an appealing choice for fencing of the MPA output, as their computational cost is $\mathcal{O}(1)$ per output value in both time and storage, and only two of such filters are needed to continuously monitor both the central tendency and the dispersion.

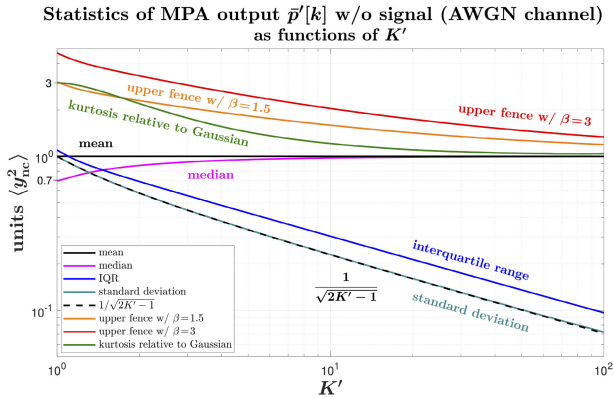


FIGURE 18. MPA filtering is computationally inexpensive way to reduce noise relative to peak signal power in leading pulse sequence.

In brief, the signal $Q_q(t)$ that is related to the given input $x(t)$ by the equation

$$\frac{d}{dt} Q_q = \mu \left[\lim_{\varepsilon \rightarrow 0} \mathcal{S}_\varepsilon(x - Q_q) + 2q - 1 \right], \quad (22)$$

where μ is the rate parameter and $0 < q < 1$ is the quantile parameter, can be used to approximate (“track”) the q -th quantile of $x(t)$. In (22), the comparator function $\mathcal{S}_\varepsilon(x)$ can be any continuous function such that $\mathcal{S}_\varepsilon(x) = \text{sgn}(x)$ for $|x| \gg \varepsilon$, and $\mathcal{S}_\varepsilon(x)$ changes monotonically from “−1” to “1” so that most of this change occurs over the range $[-\varepsilon, \varepsilon]$. In particular, as discussed in detail in [14], for a continuous stationary signal $x(t)$ with a constant mean and a positive IQR, the outputs $Q_{[1]}(t)$ (for $q = 1/4$) and $Q_{[3]}(t)$ (for $q = 3/4$) of QTFs with a sufficiently small rate parameter μ will approximate the 1st and the 3rd quartiles, respectively, of the signal obtained in a moving boxcar time window with the width ΔT of order $2 \times \text{IQR}/\mu \gg \langle f \rangle^{-1}$, where $\langle f \rangle$ is the average crossing rate of $x(t)$ with the 1st and the 3rd quartiles of $x(t)$. In numerical implementations, the sign function can be used as the comparator function, as illustrated in Appendix C.

Hence the robust fence $\alpha'[k]$ can be obtained as the upper Tukey’s fence [22] constructed as a linear combination of the outputs $Q_{[1]}[k]$ and $Q_{[3]}[k]$ for the MPA input:

$$\alpha'[k] = Q_{[3]}[k] + \beta (Q_{[3]}[k] - Q_{[1]}[k]), \quad (23)$$

where β is a positive scaling parameter of order unity (e.g., $\beta \in [1, 3]$). The overall behavior of such QTF fencing for a stationary constant-mean signal with a given IQR would be similar to the fencing with the “exact” quartile filters in a moving boxcar window of length $\Delta T = 2 \times \text{IQR}/\mu$. However, numerical computations of such exact quartiles become prohibitively expensive for $F_s \Delta T \gg 1$. This would be the case when the time interval ΔT spans multiple interpulse intervals. On the other hand, the computational cost of QTF remains $\mathcal{O}(1)$ per output value in both time and storage, and the contribution of the QTF fencing in the overall computational burden would be negligible.

Note that (see Fig. 18) increasing K' in the MPA filtering quickly reduces the kurtosis of the noise, and thus the rate of

noise outliers, down to that of Gaussian distribution. Thus, for larger β in (23) and/or K' , the rate of occurrence \mathcal{R} of the $\bar{p}'[k]$ peaks that are above the fence and are due to the noise can be made negligibly small. At the same time, the rate of peaks above $\alpha'[k]$ in the leading pulse sequence approaches F_s/N'_p when the signal becomes sufficiently strong. Therefore, as illustrated in Fig. 19, the increase in rate \mathcal{R} of peaks that are above $\alpha'[k]$ in the MPA output can reliably indicate onset of the leading pulse sequence even at low SNRs. Hence, for detection of the leading sequence in this prototype algorithm, we require a relatively large group of $\bar{p}'[k]$ peaks with IpIs N'_p to protrude above $\alpha'[k]$.

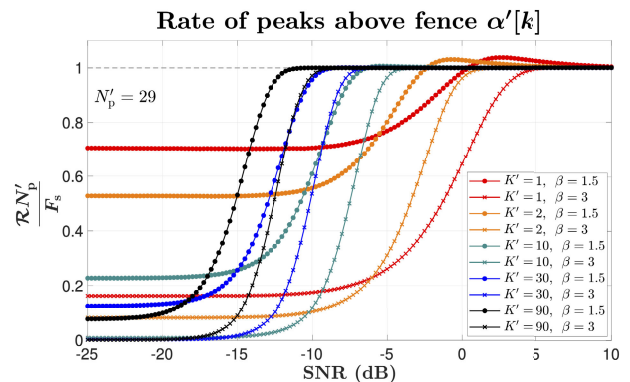


FIGURE 19. Increase in rate \mathcal{R} of peaks that are above $\alpha'[k]$ in MPA output can reliably indicate onset of leading pulse sequence.

D. PACKET DETECTION COMBINED WITH CFO MEASUREMENTS AT LOW SNR

One of the main objectives in the development of the packet detection and CFO measurement algorithm is that its performance is appropriately “matched” with that of the synchronization and decoding of the payload. Indeed, there is limited utility in the correct detection of the packets, and obtaining the CFO, if the subsequent payload cannot be reliably synchronized and decoded. Similarly, it would be wasteful to “miss” high-quality payloads by failing to correctly detect the leading pulse sequence and obtain the CFO. Let us now examine if, and how, such a match can be achieved in practice, and what are the available tradeoffs. Then the particular algorithm parameters can be appropriately chosen.

First, Figs. 20 and 21 illustrate the waveforms and the performance of the algorithm described in Section VII-B (for the same algorithm parameters and the CFO values) in the presence of noise with the same level as in Fig. 5 ($\Gamma = -10$ dB). In this example, even in the presence of dominant noise, we still correctly detect the arrival of the M-ASPM packets and accurately determine their CFOs.

In Fig. 20, we also show the MTF output $q[k]$, which is effectively the same for both $y_{nc}^2[k]$ and $y_{nc}^2[k]$. Thus $q[k]$ can be “handed over” to the execution of the basic algorithm, and used for the initialization of the MPA output $\bar{p}[k]$ (e.g., by setting its initial N'_p values to $1.4 q[k]$).

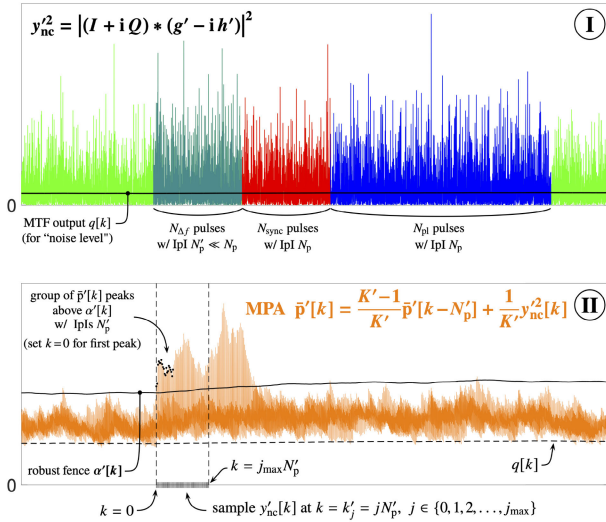


FIGURE 20. Example of waveforms in prototype algorithm with same noise level as in Fig. 5 ($\Gamma = -10$ dB).

Further, Fig. 22 illustrates the general relation between the fence scaling parameter β and the value of K' in the MPA filtering, in their impact on the probability P_d of packet detection with correct (to within $\pm F_s/(58j_{\max})$) CFO measurement. Specifically, for the simulated results shown in this figure, we maintain $N_p = 29$, $N_{\Delta f} = 6K'$, $j_{\max} = 4K'$, and use $2K'$ as the length of the group of $\tilde{p}'[k]$ peaks above $\alpha'[k]$. The values of the CFO are generated, with uniform probability, in the interval $[-F_s/60, F_s/60]$.

As can be seen in Fig. 18, the noise variance is inversely proportional to K' for large K' . If we define the “detection sensitivity” for a given P_d as $1/\Gamma'(P_d)$, then, for a given β , this sensitivity generally increases roughly proportional to $\sqrt{K'}$, $\Gamma'(P_d) \propto 1/\sqrt{K'}$. In addition, for larger values of K' we can also “tighten” the fence $\alpha'[k]$ (i.e., reduce β), further enhancing the sensitivity. Such supplementary sensitivity increase from 4.4 dB to 6 dB is emphasized in Fig. 22 for $\{\beta, K'\} = \{2, 10\}$, $\{2, 90\}$, and $\{1.2, 90\}$. Consequently, $\Gamma'(P_d) \propto 1/K'^{\kappa}$, $1/2 < \kappa < 1$, where we can use $\kappa \approx 2/3$ for rough estimates.

Further, note that the pulse duty cycle in the leading sequence is unity, and the power of the leading sequence is always higher than the average power of the payload (hence $\Gamma' = \Gamma/D$ in Fig. 22). This is a simple consequence of $\text{PAPR} = 1/D$ for the M-ASPM TX signal. Thus, for a given payload SNR Γ , the desired detection probability P_d can be achieved by adjusting either K' or the payload pulse duty cycle D , or both. In particular, for a given Γ in Fig. 22, approximately the same detection probability is achieved for $\{D, K'\} = \{D_0, 90\}$ and $\{D, K'\} = \{0.25 D_0, 10\}$:

$$P_d(\Gamma; D = D_0, K' = 90) \approx P_d(\Gamma; D = 0.25 D_0, K' = 10). \quad (24)$$

According to (24), for increasing the detection probability, a 4-fold reduction in the payload pulse duty cycle is roughly equivalent to a 9-fold extension of the leading pulse sequence.

Therefore, as a rule of thumb, and subject to the particular technical and/or regulatory constraints (such as, for example, the maximum available or allowed peak transmit power), we may want to use the smallest possible pulse duty cycle for the timing and payload sequences. A smaller duty cycle allows us to use a shorter leading sequence, offering clear benefits in terms of the latency of the detection (which is proportional to K'), and reducing the computational cost of both the detection and the subsequent signal processing for synchronization and payload decoding.

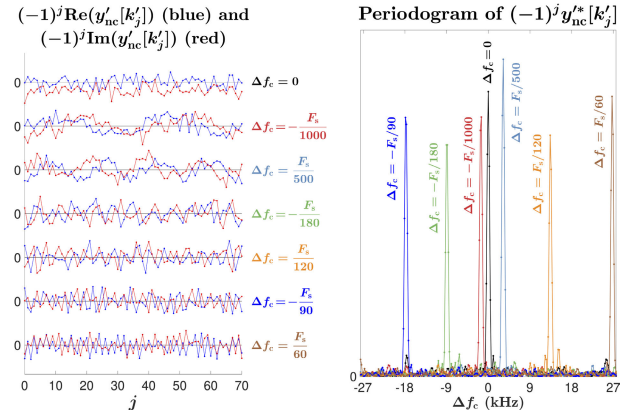


FIGURE 21. Measured CFOs for noisy waveforms ($k'_j = jN_p$).

Packet detection w/ correct CFO measurement




FIGURE 22. Probability of packet detection with correct CFO measurement.

Also, for a given average transmit power of the payload, the total energy of the timing pulse sequence and the payload is not affected by the pulse duty cycle. On the other hand, the energy of the leading sequence with the same peak amplitude is proportional to K'/D , $\mathcal{E}_{\Delta f} \propto K'/D$. Consequently, for a given detection probability, both the latency and the energy of the leading sequence can be related to the pulse duty cycle as

$$K' \propto D^{\frac{1}{\kappa}} \approx D^{1.5}, \quad \mathcal{E}_{\Delta f} \propto D^{\frac{1-\kappa}{\kappa}} \approx \sqrt{D}, \quad (25)$$

where we use $\kappa = 2/3$ for rough estimates. This is illustrated in Fig. 23. Note that $\Gamma = -13$ dB used in this example approximately corresponds to $\text{BER} = 10^{-4}$ in noncoherent LoRa with SF = 9 (see Fig. 3).

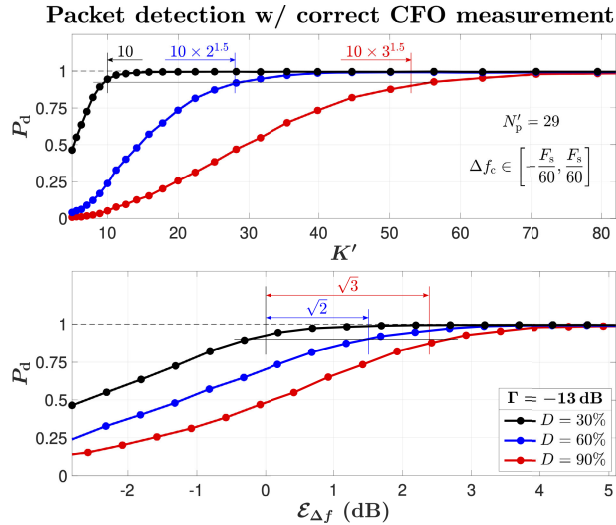


FIGURE 23. Impact of pulse duty cycle on detection latency and transmit energy of leading sequence.

VIII. DISCUSSION

In the development of the implementation framework presented in this paper, we have implicitly assumed an M-ASPM network configuration where the packets are transmitted by multiple uplink nodes, and received by a single stationary gateway.

For a single packet, the digital modulating signal $z[k]$ generated in a TX can be expressed, for example, as

$$z[k] = \sqrt{L'} \sum_i \hat{\xi}'[k - k'_i] (-1)^i + \sqrt{L} \sum_j \hat{\xi}[k - k_j] (-1)^j, \tag{26}$$

where $\hat{\xi}'[k]$ is the PSF for the leading pulse sequence, $\hat{\xi}[k]$ is the PSF for the timing and the payload sequences, and the information is encoded in the pulse-position offsets of the k_j values corresponding to the payload. (In (26) all k'_i values are smaller than any of k_j .) No other signal processing, beyond summing such time-delayed waveforms, is performed in the TX. After the D/A conversion ($z[k] \rightarrow z(t)$), the real and imaginary parts of $z(t)$ are used for quadrature amplitude modulation of a carrier with frequency f'_c , amplified, and transmitted. The alternating signs of the waveforms in (26) simply ensure that $z(t)$ is a zero-mean signal even when the individual PSFs are not.

For better TX efficiency, we require that the pulses in (26) do not overlap, and a packet consists of constant-envelope pulses with the same amplitude. Then, if the “idle” (i.e., for the zero-amplitude intervals between pulses) power consumption during the transmission of such a signal is negligible, the efficiency of this transmission will be effectively the same as the efficiency of transmitting a continuous constant-envelope waveform.

In the RX (with the LO frequency f_c), we assume that the bandwidths of the bandpass and the lowpass/anti-aliasing filters are sufficiently large to accommodate the CFO $\Delta f_c = f'_c - f_c$ in the desired range without attenuation.

Since a single gateway can receive and process packets from a large number (e.g., thousands) of TXs with different CFOs, any hardware adjustments in the RX (e.g., in the LO frequency or sampling time offsets) to compensate for the CFO differences would be undesirable. Such hardware changes are especially unwelcome for simultaneous processing of colliding packets in different PSF channels. Consequently, the CFO compensation is performed entirely in the RX software.

We use the leading pulse sequence in a packet to detect the packet, and to obtain the CFO value for this packet.

In general, for maximum detection sensitivity, we want the length $L' = N'_p$ of the PSF $\hat{\xi}'[k]$ to be as large as possible, while still providing the CFO measurements in the desired range, for the given bandwidth and the carrier frequency. For example, for the ± 30 ppm CFO and the bandwidth $B = 500$ kHz, we can use $L' = 29$ for $f_c = 915$ MHz, and increase it to $L' = 61$ for $f_c = 434$ MHz.

The other parameters of the prototype algorithm described in Section VII should be chosen to appropriately match its performance with that of the synchronization and decoding of the payload. That is, the values of K' and β (as well as $N_{\Delta f}$, j_{max} , and the length of the group of $\hat{p}'[k]$ peaks above $\alpha'[k]$) would be based on the values of the IpI N_p , the length L of the PSF $\hat{\xi}[k]$ (or, equivalently, the pulse duty cycle D), and the number M of the pulse-position states in the encoding.

Once the CFO Δf for the detected packet is measured, we can modify the matched filter $\zeta[k]$ in the RX according to (16) (i.e., $\zeta[k] \rightarrow \zeta_{\Delta}[k]$), and obtain the value of δk_1 from (17). We can then use $\zeta_{\Delta}[k]$ in the basic algorithm of Section III for the synchronization (as well as for measuring the signal quality), while implementing the cumulative STO compensation described in Section VI for the payload sampling.

A. SIMULATED PERFORMANCE OF FULL ALGORITHM

Let us now illustrate performance of a full prototype algorithm with the example shown in Fig. 24. In this illustration, we use 64-ASPM with the processing gain 32.3 dB which, for negligible CFO and STO, provides the uncoded BER = 10^{-4} at the SNR $\Gamma \approx -15.6$ dB. Note that this BER performance effectively matches LoRa with SF = 10, as shown by the solid and dashed black lines in Fig. 24 (also see Fig. 3).

In the simulations, we use 120 pulses (i.e., 90 bytes) for the payload, and 9 pulses for the timing sequence. For $B = 500$ kHz bandwidth, the time duration of the TX packet shown at the top of Fig. 24 is about 222 ms. For the timing and the payload sequences we use a relatively short PSF with $L = 600$, and thus the pulse duty cycle is $D \approx 22\%$. Then, with $f_c = 915$ MHz carrier frequency (that is, $f_c/B = 915 \text{ MHz}/500 \text{ kHz} = 1,830$), $\Delta f_c \approx \pm F_s/L$ for the ± 3 ppm CFO, and $\Delta f_c \approx \pm 10 F_s/L$ for the ± 30 ppm CFO. As one can see in the left panel of Fig. 9 (see Section V), the reduction

in the peak magnitude of the received pulses is less than 1 dB for ± 3 ppm, and more than 6 dB for ± 30 ppm.

Further, the values of δk_1 are approximately $\mp 8.2 \times 10^{-3}$ and ∓ 0.082 , for the ± 3 ppm and ± 30 ppm CFOs, respectively. In the first case the magnitude of the difference in cumulative CTO remains below approximately ∓ 1 for the duration of the packet, while in the second case the deterioration in the RX signal quality is about 37% faster than shown in the upper panels of Figs. 11, 12, and 14 (see Section VI).

For the leading sequence we employ 180 pulses with $L' = N'_p = 29$. As such, with these parameter values, the leading sequence spans less than 2 IpIs N_p , and its time duration is about 3.3 ms ($\approx 1.5\%$ of the packet length). Further, we use $K' = 60$, $j_{\max} = 120$, and $\beta = 2$ in the detection algorithm.

In the simulations, the CFO values are generated with uniform probability in three different intervals: ± 3 ppm, ± 10 ppm, and ± 30 ppm. For each CFO range, we generate 60,000 packets with initial STO values randomly chosen in the ± 0.5 interval.

In the right axis of Fig. 24, we plot the “raw” packet detection probability. That is, we do not check if the CFO value was obtained correctly. Note that this raw probability is different from the joint probabilities of detection with correct CFO measurements shown in Figs. 22 and 23.

After a packet is detected and the measured CFO value is obtained, we modify the RX matched filter and attempt to acquire the synchronization and decode the payload. To account for the impact of false-negative synchronizations, if, after the packet detection, we do not obtain the synchronization during the combined duration of the leading and the timing sequences, we assign a random value (within this interval) for the synchronization time. Then, for the detected packets, we decode the payload and obtain the BER values. For the payload sampling, we use a simple STO compensation to within ± 1 according to (18).

For comparison, in Fig. 24 we also provide the BER values when using unmodified RX matched filter and sampling without cumulative STO compensation.

B. MULTIPLE PSF CHANNELS

The detection sensitivity of the prototype algorithm described in Section VII is high enough to enable the use of relatively sensitive M-ASPM configurations with large IpIs, for example, matching LoRa with SF = 10 in terms of the path loss and the range (see Fig. 24). At the same time, such large-IpI configurations offer proportionately large SIR margins. Then we can employ several different PSF channels, respectively increasing the network capacity (e.g., to several times that of LoRa with SF = 10) and/or its effective range [2], [3]. Within the framework of this paper, such multiple PSF channels can be enabled, for example, in the following manner.

First, we can construct two different leading pulse sequences using “flip” PSFs $\hat{\zeta}'_1[k]$ and $\hat{\zeta}'_2[k]$, such that the

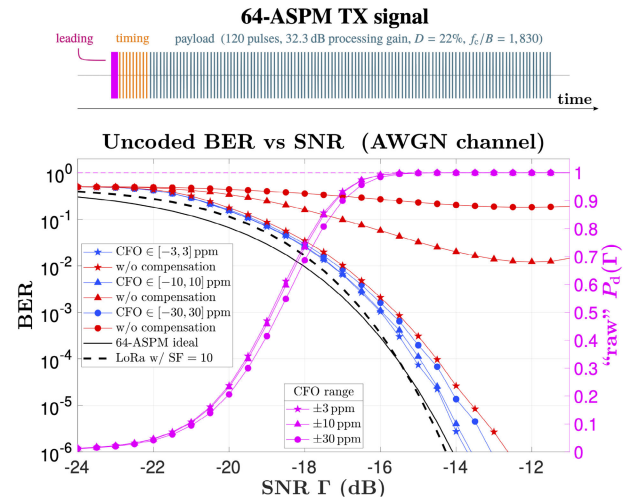


FIGURE 24. Example of simulated BER performance of full algorithm for different ranges of CFO values. Right axis: “Raw” probability of packet detection with leading sequence.

2nd PSF is the complex conjugate of the matched filter for the 1st PSF, $\hat{\zeta}'_2[k] = \hat{\zeta}'_1*[k]$ [2], [3]. Consequently, these two different leading sequences can be used for two different PSF channels (i.e., with different PSFs $\hat{\zeta}_1[k]$ and $\hat{\zeta}_2[k]$).

In addition (or as an alternative), different M-ASPM packets can share the same leading and timing sequences, and use the same PSFs for short (e.g., shorter than the timing sequence) headers that include the PSF channel information. Then much longer payloads, that use different PSFs specified in the respective headers, can be used for the rest of the packets.

C. FILTERING AND SAMPLING

The M-ASPM presentation in Section II, as well as the description of the basic algorithm in Section III, suggest that we need to perform continuous FIR filtering at full sampling rate in order to detect and synchronize the incoming signal. Since the length L of the “main” matched filter $\zeta[k]$ in the receiver may be quite large (e.g., hundreds or thousands of samples), such continuous filtering would be a computationally expensive operation, especially wasteful for low packet rates.

Favorably, such continuous filtering only needs to be performed for detection of the leading pulse sequence, specifically, for detection of the $\bar{p}'[k]$ onset. This uses significantly shorter FIR filters, of the length $L' = N'_p \ll N_p$, and is much less computationally intensive. Further, the MPA filter for obtaining $\bar{p}'[k]$ is effectively just a 1st order IIR filter (albeit with $\mathcal{O}(N'_p)$ per output value in the storage cost), and the quantile tracking filters used for constructing the robust fence $\alpha'[k]$ are also numerically inexpensive.

Only after the packet has been detected and the CFO has been measured, we begin filtering with the modified matched filter $\zeta_\Delta[k]$ of length L , and obtaining the MPA output $\bar{p}[k]$. Initially, until the onset of $\bar{p}[k]$ is detected (e.g., as two consecutive peaks of $\bar{p}[k]$ that are above the fence $\alpha[k]$ and

are IpI N_p apart), we would need to obtain the filtered output at full sampling rate. However, after detecting the onset of $\bar{p}[k]$, we no longer need to continue the MPA filtering, and only need to obtain a decimated output of $y_{nc}[k]$. Initially, until the last peak in the timing sequence is detected, we need to acquire $M + 1$ samples per IpI N_p . For the payload, we obtain M samples per IpI.

IX. CONCLUSION

Among appealing features of M-ASPM is its extensive versatility in tailoring multiple PHY parameters to reconcile often conflicting LPWAN technical concerns. Since carrier synchronization at low SNRs is not trivial, a major such concern is maintaining the M-ASPM signal quality, for a wide range of data rates and payload sizes, under rather significant (e.g., ± 30 ppm) carrier frequency offsets. Further, it is desirable to achieved this without any feedback communications between the TX and RX, any hardware or software changes in the TX, and any hardware adjustments in the RX (e.g., in the LO frequency or sampling time offsets). In this paper, we outline a practical framework for combining the detection, synchronization, and decoding of M-ASPM packets in the presence of such CFOs, and suggest a prototype algorithm for its implementation.

In our descriptions and illustrations of the algorithm steps, we aim to provide sufficient reasoning and details for (i) enabling a working SDR prototype of an M-ASPM system, and (ii) facilitating the subsequent adaptations of the algorithm, under given technical constraints, to specific practical challenges. These challenges can include, for example, significant delay spreads, external technogenic interference, and co-PSF and inter-PSF collisions.

APPENDIX A ACRONYMS

ACF: autocorrelation function; A/D: Analog-to-Digital; ASPM: Aggregate Spread Pulse Modulation; AWGN: Additive White Gaussian Noise; BER: Bit Error Rate; CFO: Carrier Frequency Offset; D/A: Digital-to-Analog; FIR: Finite Impulse Response; IIR: Infinite Impulse Response; IQR: Interquartile Range; IpI: Interpulse Interval; LoRa: Long Range (modulation technique for LPWANs based on chirp spread spectrum); LO: Local Oscillator; LPWAN: Low-Power Wide Area Network; M-ASPM: M-ary ASPM; mph: miles per hour; MPA: Modulo Power Averaging; MTF: Median Tracking Filter; PAPR: Peak-to-Average Power Ratio; PHY: physical layer; ppm: parts per million; PSD: Power Spectral Density; PSF: Pulse Shaping Filter; QTF: Quantile Tracking Filter; RC: Raised-Cosine; RX: receiver; SDR: Software-Defined Radio; SF: Spreading Factor (for LoRa); SFO: Sampling Frequency Offset; STO: Sampling Time Offset; SIR: Signal-to-Interference Ratio; SINR: Signal-to-Interference-plus-Noise Ratio; SNR: Signal-to-Noise Ratio; SQI: Signal Quality Indicator; STO: Sampling Time Offset; TBP: Time-Bandwidth Product; ToA: Time-on-Air; TX: transmitter.

APPENDIX B

COMMENTS ON NOTATIONS

Whenever a particular notation is introduced in the paper, it is immediately defined. Some notations are confined to the specific sections. The notations that are used multiple times throughout the paper include:

α	upper fence/threshold
B	bandwidth
β	RC roll-off factor or scaling parameter in QTF fencing
D	pulse duty cycle
Δf_c	CFO value
η	spectral efficiency
f_b	bit rate
f_c	carrier/LO frequency
F_s	sample rate
Γ	SNR
k	sample index (in digital signal representations)
K	MPA filter parameter
L	PSF length
λ	signal quality indicator
M	number of states in M-ary encoding
μ	rate parameter in QTFs
N_p	average interpulse interval
N_{pl}	number of pulses in payload
\mathcal{N}_s	oversampling factor
\bar{p}	output of MPA filter
P_b	bit error probability
P_d	probability of packet detection
q	MTF output
y_{nc}	received pulse train (noncoherent detection)

In the mathematical notations we reserve the letters “ ζ ”, “ g ”, and “ h ” for pulse shaping filters, with g and h being the real and imaginary parts of ζ . For example, we denote the finite impulse response of a PSF applied to a designed pulse train as $\hat{\zeta}[k]$, where k is the sample index. Since in this paper we assume the single-sideband M-ASPM, the PSF components g and h are related through the Hilbert transform, e.g., $h(t) = \pm H(g(t))$ (in analog domain) or $h[k] = \pm H\{g[k]\}$ (in digital representation).

Further, we find it convenient to use the “hat” operator for $\hat{\zeta}[k]$, $\hat{g}[k]$, and $\hat{h}[k]$ to distinguish them from their respective matched filters $\zeta[k] = \hat{\zeta}^*[-k]$, $g[k] = \hat{g}[-k]$, and $h[k] = -\hat{h}[-k]$. We also use the hat symbol in Sections II and VII to denote the designed pulse train $\hat{x}[k]$, as opposed to the shaped train $z[k]$ obtained by applying $\hat{\zeta}[k]$ to the designed pulse sequence.

To distinguish between the respective quantities for LoRa and M-ASPM, those for LoRa are marked by overhead tildes.

We use primes to denote the quantities related specifically to the leading pulse sequence, as opposed to the same quantities for the timing and the payload sequences. We also use primes to distinguish the carrier and the sampling frequencies in the TX from those in the RX.

APPENDIX C

NUMERICAL IMPLEMENTATION OF QTFS

Since outputs of analog QTFS are piecewise-linear signals consisting of alternating segments with positive and negative slopes, care should be taken in finite difference implementations of QTFS to avoid the “overshoots” around the crossings of $Q_q(t)$ with $x(t)$. In particular, when $x[k] - Q_q[k-1]$ is outside the interval $h\mu [2(q-1), 2q]$, where h is the time step, one may set $Q_q[k] = x[k]$, as illustrated in the MATLAB function below:

```

%-----
function y = QTFS(x, dt, mu, q)
%-----
lx = length(x); lq = length(q);
q = q(:); y = zeros(lx, lq); gamma = mu*dt;
y(1, :) = x(1) * ones(1, lq);
%-----
for i = 2:lx
    dX = x(i) * ones(1, lq) - y(i-1, :);
    for j = 1:lq
        if dX(j) > 2*gamma*(q(j)-1) & dX(j) < 2*gamma*q(j)
            y(i, j) = x(i);
        else
            y(i, j) = y(i-1, j) + gamma * (sign(dX(j)) + 2*q(j) - 1);
        end
    end
end
return
%-----

```

REFERENCES

- [1] A. V. Nikitin and R. L. Davidchack, “M-ary aggregate spread pulse modulation in LPWANs for IoT applications,” in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Madrid, Spain, Dec. 2021, pp. 1–7, doi: [10.1109/GLOBECOM46510.2021.9685345](https://doi.org/10.1109/GLOBECOM46510.2021.9685345).
- [2] A. V. Nikitin and R. L. Davidchack, “M-ary aggregate spread pulse modulation for robust and scalable low-power wireless networks,” *IEEE Access*, vol. 10, pp. 96652–96671, 2022, doi: [10.1109/ACCESS.2022.3205009](https://doi.org/10.1109/ACCESS.2022.3205009).
- [3] A. V. Nikitin and R. L. Davidchack, “M-ary aggregate spread pulse modulation with pulse-shaping power control for highly scalable LPWANs,” *IEEE Access*, vol. 11, pp. 25642–25661, 2023, doi: [10.1109/ACCESS.2023.3256888](https://doi.org/10.1109/ACCESS.2023.3256888).
- [4] L. Vangelista, “Frequency shift chirp modulation: The LoRa modulation,” *IEEE Signal Process. Lett.*, vol. 24, no. 12, pp. 1818–1821, Dec. 2017, doi: [10.1109/LSP.2017.2762960](https://doi.org/10.1109/LSP.2017.2762960).
- [5] G. Baruffa, L. Rugini, L. Germani, and F. Frescura, “Error probability performance of chirp modulation in uncoded and coded LoRa systems,” *Digit. Signal Process.*, vol. 106, Nov. 2020, Art. no. 102828, doi: [10.1016/j.dsp.2020.102828](https://doi.org/10.1016/j.dsp.2020.102828).
- [6] A. V. Nikitin, “Communications method and apparatus,” U.S. Patent 11 671 290, Jun. 6, 2023.
- [7] D. E. Knuth, “Two notes on notation,” *Amer. Math. Monthly*, vol. 99, no. 5, pp. 403–422, May 1992, doi: [10.2307/2325085](https://doi.org/10.2307/2325085).
- [8] A. W. Doerry, “Generating precision nonlinear FM chirp waveforms,” in *Radar Sensor Technology XI*, vol. 6547, J. L. Kurtz and R. J. Tan, Eds. Bellingham, WA, USA: SPIE, 2007, Art. no. 65470D, doi: [10.1117/12.717796](https://doi.org/10.1117/12.717796).
- [9] R. N. Bracewell, *The Fourier Transform and Its Applications*, 3rd ed. New York, NY, USA: McGraw-Hill, 2000.
- [10] G. Todoran, R. Holonec, and C. Iakab, “Discrete Hilbert transform. Numeric algorithms,” *Acta Electroteh.*, vol. 49, no. 4, pp. 485–490, 2008.
- [11] D. Gabor, “Theory of communication,” *J. Inst. Electr. Eng.*, vol. 93, no. 26, pp. 429–457, 1946, doi: [10.1049/ji-3-2.1946.0074](https://doi.org/10.1049/ji-3-2.1946.0074).
- [12] M. Vetterli and J. Kovačević, *Wavelets and Subband Coding*. Upper Saddle River, NJ, USA: Prentice-Hall, 1995.
- [13] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms, and Applications*, 4th ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2006.
- [14] A. V. Nikitin and R. L. Davidchack, “Quantile tracking filters for robust fencing in intermittently nonlinear filtering,” 2019, *arXiv:1911.00736*, doi: [10.48550/arXiv.1911.00736](https://doi.org/10.48550/arXiv.1911.00736).
- [15] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions*. New York, NY, USA: Dover, 1972.
- [16] H. A. David and H. N. Nagaraja, *Order Statistics* (Wiley Series in Probability and Statistics), 3rd ed. Hoboken, NJ, USA: Wiley, 2004.
- [17] T. S. Rappaport, *Wireless Communications: Principles and Practice* (Prentice Hall Communications Engineering and Emerging Technologies Series), 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2002.
- [18] A. A. Doroshkin, A. M. Zadorozhny, O. N. Kus, V. Yu. Prokopyev, and Y. M. Prokopyev, “Experimental study of LoRa modulation immunity to Doppler effect in CubeSat radio communications,” *IEEE Access*, vol. 7, pp. 75721–75731, 2019, doi: [10.1109/ACCESS.2019.2919274](https://doi.org/10.1109/ACCESS.2019.2919274).
- [19] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd ed. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [20] A. V. Nikitin and R. L. Davidchack, “Analog-domain mitigation of outlier noise in the process of analog-to-digital conversion,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, Kansas City, MO, USA, May 2018, pp. 1–7, doi: [10.1109/ICC.2018.8423032](https://doi.org/10.1109/ICC.2018.8423032).
- [21] A. V. Nikitin and R. L. Davidchack, “Pulsed waveforms and intermittently nonlinear filtering in synthesis of low-SNR and covert communications,” *IEEE Access*, vol. 8, pp. 173250–173266, 2020, doi: [10.1109/ACCESS.2020.3024978](https://doi.org/10.1109/ACCESS.2020.3024978).
- [22] J. W. Tukey, *Exploratory Data Analysis*. Reading, MA, USA: Addison-Wesley, 1977.



ALEXEI V. NIKITIN (Member, IEEE) is a Co-Founder and Chief Science Officer of the Kansas-based Nonlinear LLC. He initiated his undergraduate and graduate studies in physics, chemistry, and engineering in the former USSR at Novosibirsk State University in Novosibirsk and Karpov Institute of Physical Chemistry in Moscow. After receiving a Ph.D. degree in physics from the University of Kansas in 1998. He led R&D work focused on methods and tools in nonlinear signal processing at several startup companies, some of which were subsequently acquired, specializing in applications in communications, power electronics, navigation, geophysical sciences, neurology, and biometrics. He is named as a first or sole inventor on over 30 issued US patents.



RUSLAN L. DAVIDCHACK is Professor of Mathematical Modelling and Computation with the School of Computing and Mathematical Sciences, University of Leicester, U.K. His research interest is in developing computational methods and tools for applications in molecular simulations, nonlinear dynamics and signal processing.

Dr. Davidchack received his undergraduate degree in theoretical condensed matter physics from Lviv University, Ukraine, and the Ph.D. degree in computational statistical physics from the University of Kansas. He has developed several unique computational approaches and made significant contributions to a variety of fields in computational chemistry, signal processing, mathematics, statistical physics, and biochemistry. He has an extensive track record of interdisciplinary collaborations which span academic institutions (e.g. Georgia Tech; UC Berkeley; Oxford, U.K.; EPFL, Switzerland; NTNU, Norway; KTH, Sweden), national laboratories in the USA (e.g. AFRL, Oak Ridge, Sandia), and U.K. (Daresbury, Institute of Pharmaceutical Innovation), and industrial partners (Tata Steel, Rolls-Royce, GlaxoSmithKline, Procter & Gamble). He has given invited talks at many international conferences in the USA, U.K., France, Germany, South Korea, China, and Ukraine.

•••