IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received July 17, 2019, accepted August 6, 2019, date of publication August 27, 2019, date of current version September 11, 2019.

Digital Object Identifier 10.1109/ACCESS.2019.2937958

Fallthrough Correlation Techniques for Arbitrary-Phase Spread Spectrum Waveforms

MICHAEL FLETCHER, (Student Member, IEEE), ALAN MICHAELS, (Senior Member, IEEE), AND DEVIN RIDGE, (Member, IEEE)

Hume Center for National Security and Technology, Virginia Polytechnic Institute and State University, Blacksburg, VA 24060, USA

Corresponding author: Michael Fletcher (mjf@vt.edu)

ABSTRACT The use of practically non-repeating spreading codes to generate sequence-based spread spectrum waveforms is a strong method to improve transmission security, by limiting an observer's opportunity to cross-correlate snapshots of the signal into a coherent gain. Such time-varying codes, particularly when used to define multi-bit resolution arbitrary-phase waveforms, present significant challenges to the intended receiver, who must synchronize acquisition processing to match the time-varying code each time it changes. This paper presents a series of options for optimizing the traditional brute-force matched-filter preamble correlator for burst-mode arbitrary-phase spread spectrum signals, achieving significant computational gains and flexibility, backed by measurable results from hardware prototypes built on an Intel Arria 10 Field Programmable Gate Array (FPGA). The most promising of which requires no embedded multipliers and reduces the total hardware logic by more than 76%. Extensions of the core fallthrough correlator techniques are considered to support low-power asynchronous reception, underlay-based physical layer firewall functions, and Receiver-Assigned Code Division Multiple Access (RA-CDMA) protocols in Internet of Things (IoT)-caliber devices.

INDEX TERMS Spread spectrum, Internet of Things (IoT), chaotic communications, signal detection, correlation, FPGA.

I. INTRODUCTION

The design of burst-mode communication systems presents additional challenges over that of a standard continuous data link, in particular due to the need to re-acquire the signal on a burst-by-burst basis. In low-power devices, such as those suitable for Internet of Things (IoT), burst-mode waveforms traditionally employ techniques to make the acquisition preamble as easy to receive as possible, typically by embedding pilot tones [1], repeated cyclic prefixes [2], cyclic autocorrelation functions [3], soft-handoff between spreading codes [4], Barker-sequence / short preamble repetition [5], maximal-likelihood estimation [6], and/or variations of matched-filter techniques [7], [8]. Virtually all of these approaches rely on an inherent cyclostationary signal feature of the preamble bursts, facilitating blind detection and/or exploitation by an unintended receiver.

All signals considered in this paper are digital chaotic sequence-based arbitrary-phase spread spectrum waveforms with optional chip amplitude shaping, most of which use practically non-repeating spreading codes designed to eliminate cyclostationary signal content. Reception of these signals is more complicated, using methods that adapt some aspects of the matched-filter/coherent receiver processing architectures for specific waveforms and/or use cases [9]–[12]. Further, most of these techniques are computationally intensive, making them difficult to implement in a low-power device.

Starting with the traditional brute-force matched-filter correlator, this paper presents computational efficiency improvements for a generic coherent receiver architecture where the matched-filter coefficients change on a burst-to-burst basis, offering lower-power / computationlly efficient methods that achieve the same purpose. Similar analyses have evaluated the reduced-computation processing of the semi-coherent chaotic carrier shift keying (CSK) waveforms [13]. There, however, the timing and phase are effectively coherent.

The core fallthrough correlator design model is provided in Section II. Enhancements for reduced-precision correlations, optimally pruned coefficients, and variable-length operations are all considered in Section III. Measurable results from hardware prototypes built on an Intel Arria

The associate editor coordinating the review of this article and approving it for publication was Byung-Seo Kim.

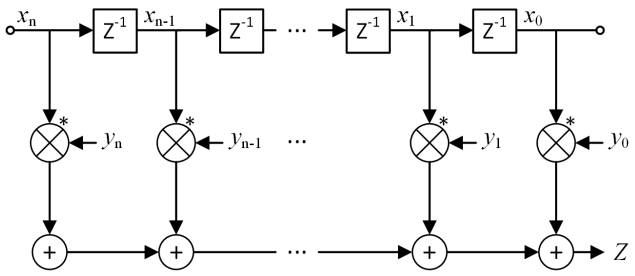

FIGURE 1. Fallthrough correlator in direct form FIR structure. The matched-filter coefficients $\{y_0, y_1, \ldots, y_n\}$ are registered for use each time epoch.

10 field-programmable gate array (FPGA) are presented in Section IV, offering simpler methods for asynchronous reception, underlay-based watermark validation [14], and receiver-assigned code-division multiple access (CDMA) operations [15] in IoT-caliber devices. Finally, conclusions can be found in Section V.

II. COMPUTATIONAL MODEL

The time-based evolution of matched-filter coefficients eliminate many of the standard methods for collapsing the digital logic structure to take advantage of a priori known cyclostationary preamble signal features, while the multi-bit resolution spreading codes used to generate the arbitrary-phase spread spectrum waveforms, such as the chaotic sequence spread spectrum (CSSS) [9], [10] or high-order PSK signaling (HOPS) [16] waveforms, increase overall complexity of the complex-conjugate multiplications (correlations) in contrast to 2-ary or 4-ary chip phase direct sequence spread spectrum (DSSS) signals [17]. To support the discussion, consider the matched-filter correlator model shown in Fig. 1, where each x is a complex-valued received signal sample and each y is a matched-filter coefficient taken from the internally generated preamble signal replica.

This model is similar to a direct form finite impulse response (FIR) filter, where a fully pipelined set of outputs are derived from incoming samples as they progress through the delay line structure. With FIR filters, significant improvements may typically be made (a) due to symmetry of wisely chosen coefficients (pre-additions), (b) elimination of sufficiently small / zero coefficients (pruning), (c) canonic signed digit (CSD) mapping of static coefficients to shift-adds [18], (d) employing computationally efficient multi-rate processing structures [19], and (e) variable control of the coefficient word widths.

Within the correlation calculation, these traditional simplifications are limited: (a) the correlation of a preamble without any cyclostationary features can not have easily exploitable symmetries, (b) coefficients may be pruned, though the correlation taps generally contribute a similar amount of energy to the composite correlation value, (c) CSD mappings may also be applied, but must be dynamically addressable with variable barrel shifters, (d) the notionally fixed sample rate hinders any multi-rate signal processing benefit, (e) and the coefficient word width in hardware will need to support the largest width that the coefficient may ever be, burst-to-burst.

In addition to these distinctions from a standard FIR filter, the logic within the fallthrough correlator must support clocking in of new coefficients on a burst-by-burst basis, so that they are in place and ready for correlation processing when the next sample arrives. Under the assumption of normalized inputs, the correlator output response Z ideally triggers based on a defined correlation peak having magnitude equal to the average chip energy times the length of the correlation. The coherent preamble signal is trivial to normalize, while the incoming received signal is variable and highly dependent on any system gains that may occur prior to the correlator. This is particularly important for spread spectrum systems, since the signal often operates at or below the ambient noise floor of the receiver and allows for power level estimates of the incoming signal and/or its multipath components based on the magnitude of the resulting correlation peak(s).

The next distinction is that of phase rotations, with particular focus on center frequency offsets. The static phase rotation may be detected from the phase offset of the correlation peak (referenced to the center of the correlation window) and subsequently corrected prior to despreading. Frequency offset, on the other hand, requires comparison of multiple sub-correlation values throughout the correlator structure, so that the phase rotations may be measured as a function of time and translated, via the known sample rate, to an instantaneous frequency offset that can be applied to the remainder of the pulse. If the frequency offset causes the correlation values to drift more than $\approx \pi/2$ radians over the duration of the preamble, then the integration process underlying the addition of the taps will begin to fail.

The final distinction of timing uncertainty due of phase noise or oscillator drift is also not supported by this FIR structure. Practical clocks (<100 ppm) will tend not to drift beyond that which is supported, and the detection of future preambles will have unique starting sample points, making only the short-term stability of individual bursts relevant.

III. FALLTHROUGH CORRELATION TECHNIQUES

The chief focus of this paper is on the computational efficiency improvements that may be made to the fallthrough correlator to achieve reasonably solid performance from a minimum amount of hardware. In particular, the allowable resource and performance trades from the hardware baseline of a brute-force design that implements a complex multiplication $z = (y_I + jy_Q)(x_I - jx_Q)$ using the three real-multiplier reduction in (1) and (2), where $z_I + jz_Q$ is a partial sum.

$$z_I = (y_I x_I + y_Q x_Q) \tag{1}$$

$$z_Q = ((y_I + y_Q)(x_I - x_Q) - y_I x_I + y_Q x_Q) \tag{2}$$

To adapt this model to an IoT-relevant context, the following series of identified improvements may be incorporated.

A. TRUNCATED COEFFICIENTS

The precision of the matched-filter coefficients may be reduced with acceptable detection loss, even for arbitrary-phase waveforms.[1] Such truncation must account for the full processing chain of the transmitted signal, including any interpolation, prior to transmission. Since each arbitrary-phase spreading chip is taken from an allowable set of discretized phase points on the unit circle [16], truncation in both the in-phase (I) component y_I and quadrature (Q) component y_Q will introduce amplitude and phase mismatch loss to the calculation. Using a bit precision $\geq$ 6 bits gives almost no performance loss, while truncation to 1-bit coefficients provides the largest computational gains, offering a hardware structure resembling the correlator of DSSS signals.

Choosing the 1-bit truncated coefficients, the correlation logic may be implemented as four negations of $\{x_I, x_Q\}$ based on sign$\{y_I, y_Q\}$ followed by two additions. Although, any quantization effects of truncation should be considered prior to processing the correlation peak for received signal power estimations. For the hardware prototypes, the overall HOPS waveform is constant envelope (i.e., $x_I^2 + x_Q^2 = 1 \forall x$), and the output response can be scaled by the reciprocal of the expected coherent signal correlation $E\left[|x_I| + |x_Q|\right] = 2 \cdot E\left[|x_I|\right] = 4/\pi$ to correct for this distortion.

B. DYNAMIC PRUNING

Upon definition of the matched-filter coefficients, if either component in I or Q does not contribute a meaningful amount of energy to the correlation, then the coefficient may be collapsed into a single real-valued or imaginary-valued correlation tap. For the spread spectrum waveforms with amplitude-varying chips, this pruning may be pursued consistent with the selective noise cancellation techniques described in [20], while for the constant envelope modulations, a parameter λ can be defined to represent the amplitudes of components to be discarded, as shown in Fig. 2.

In any scenarios where $\min\{|y_I|, |y_Q|\} < \lambda$, then $\max\{|y_I|, |y_Q|\} > \sqrt{1-\lambda^2}$, resulting in the detection performance loss shown as a function of λ in Fig. 3. The simulated loss of 0.87 dB at the median value $\lambda = \sqrt{2}/2$ allows a simplified pruning process equivalent to selecting the larger of $\{|y_I|, |y_Q|\}$ for correlations with the received signal. In other words, (1) is reduced to:[2]

$$z_I \approx \begin{cases} y_I x_I, & |y_I| \geq |y_Q| \\ y_Q x_Q, & |y_I| < |y_Q| \end{cases} \tag{3}$$

[1]The arbitrary-phase nature of these waveforms requires on the order of 2^k allowable phase words, with $k \geq 8$.

[2]By using only correlation taps that do not contain points at $|y_I| = |y_Q|$, which is easily achieved by rotating the entire set of allowed discretized points by the phase of one half LSB, a strict maximum may be achieved. In the case where the two values are equal (within the chosen comparator's precision), then the choice of which one to take forward is arbitrary.

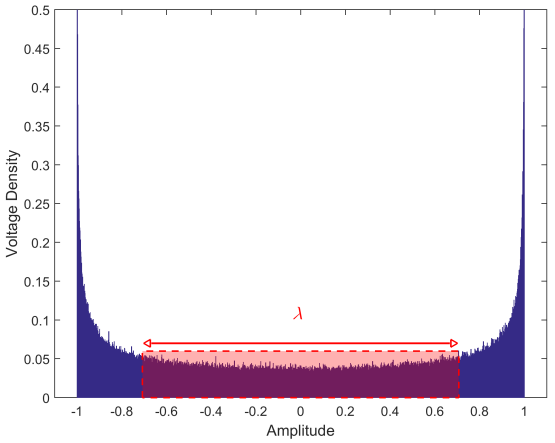

FIGURE 2. Conceptual depiction of correlation tap pruning; overlaid on voltage distribution of randomly selected points on the unit circle projected onto one axis.

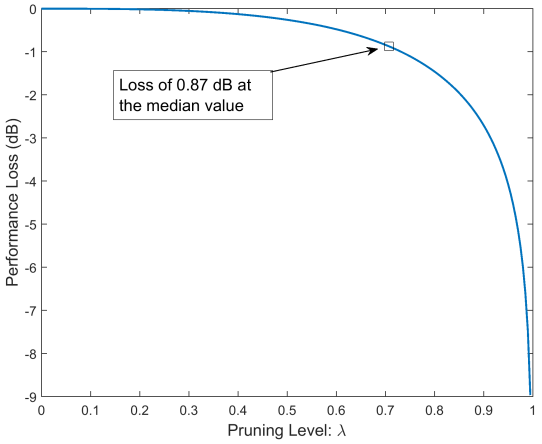

FIGURE 3. Simulated performance degradation based on choice of λ. The median value corresponds to a 0.87 dB loss (equivalently, $\lambda = \sqrt{2}/2$).

and (2) is reduced to:

$$z_Q \approx \begin{cases} -y_I x_Q, & |y_I| \geq |y_Q| \\ y_Q x_I, & |y_I| < |y_Q| \end{cases} \tag{4}$$

With 1-bit truncated $\lambda = \sqrt{2}/2$ pruned coefficients, the correlation logic may be implemented as four negations of $\{x_I, x_Q\}$ based on sign$\{y_I, y_Q\}$ followed by $\max\{|y_I|, |y_Q|\}$-induced selection of the complex-valued output. Despite the further hardware savings, pruning at the median value eliminates the amplitude mismatch by rotating the allowable taps onto the axes and reduces any phase mismatch to within $[-\pi/4, \pi/4]$, giving a degradation reduction of 3 dB over truncated coefficients without pruning. Using a smaller value for λ provides only marginal performance increases and requires dynamic placement of the adders - to allocate these adders on a burst-to-burst basis is likely to take more logic than simply provisioning all taps with the same two adder structure.

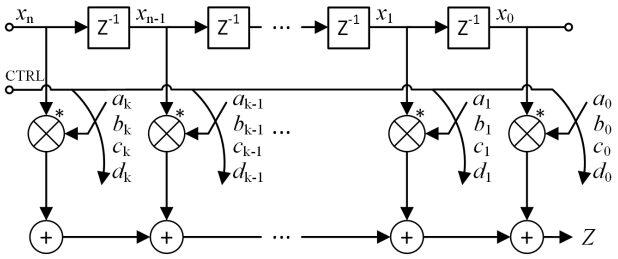

FIGURE 4. Fallthrough correlator with 4× folded correlation taps in direct form FIR structure. The control line selects one set of taps based on the pipelined detection state.

C. FOLDED CORRELATION TAPS

The sequence of matched-filter coefficients may be folded by consciously aliasing the correlation taps onto one another, achieving a significant reduction in the digital logic dedicated to the long delay lines of the received signal. While the increase in false positives would be unacceptable for lightly spread signals, the self-interference characteristics of deeply spread signals allow for minimal performance loss. Although, shorter preambles do have more difficulty in estimating phase rotations / frequency offsets, placing a practical minimum bound on the order of 2 symbols.

Consider the 4× folded correlator shown in Fig. 4, trading some additional control logic for an effectively reduced delay line length of one-fourth its original length. For the folded taps hardware prototype, the 1400 correlation taps are divided into four equal-length sets of 350 taps each. That is, $\{a_0, a_1, \ldots, a_{349}\} \equiv \{y_0, y_1, \ldots, y_{349}\}$, $\{b_0, b_1, \ldots, b_{349}\} \equiv \{y_{350}, y_{351}, \ldots, y_{699}\}$, and so on.

The control circuitry can be implemented in any number of ways. Within the context of this paper, the logic operates as follows. Each time a new incoming sample is clocked into the delay line, the control selects one of the four sub-preamble sequences to be used for correlations in that sample clock cycle. The selection is based on the pipeline decision state of previous sub-preamble detections, with reference to when the sample that just exited the delay line was the incoming sample. If a detection was triggered for the exiting sample, then the correlation taps progress to the next sequence (until another new sample is clocked in). Signal timing is acquired after four sub-preamble detections have triggered in succession.

IV. HARDWARE PROTOTYPE VALIDATION

A selection of hardware prototypes were built for reception of the arbitrary-phase HOPS spread spectrum waveform [16] and implemented on an Intel Arria 10 SoC FPGA, including: (1) a brute-force[3] matched-filter model, (2) a 1-bit truncated coefficients model, (3) a truncated coefficients model with $\lambda = \sqrt{2}/2$ pruning, and (4) a truncated pruned coefficients model with 4× folded correlation taps. The HOPS signals are

[3]The hardware prototype HOPS system employs an 8 symbol preamble and 175 chip spread ratio. Since the Arria 10 FPGA is limited to 3374 multipliers, the $3 \cdot 8 \cdot 175 = 4200$ multiply operations required by the brute-force design are clocked at a higher rate to fit on the device.

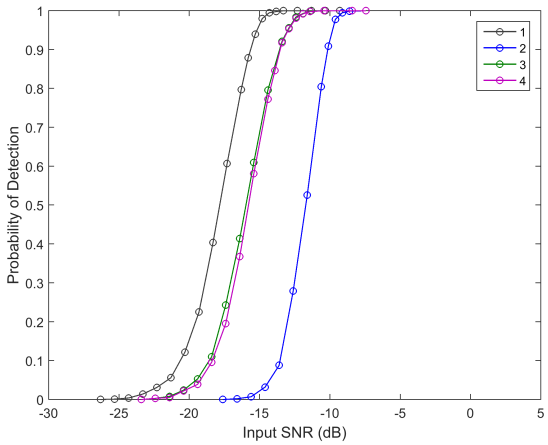

FIGURE 5. Comparative preamble detection performance for each of the fallthrough correlator design variants.

constructed in hardware using digital chaos-based spreading codes taken from an arbitrary uniform distribution of 2^k equally-spaced phase words. Given the hardware prototypes employ $k = 8$, similar results can be expected for any digital chaotic sequence-based spread spectrum waveform ($k \geq 8$.).

Of primary interest is the hardware reductions achieved by the computational enhancements, the comparative utilization numbers in Table 1 were taken from the relevant Quartus fitter reports, and focus on the use of adaptive logic modules (ALMs), combinational adaptive look-up tables (ALUTs), dedicated registers, and digital signal processing (DSP) blocks. The most significant reduction is the elimination of hardware multipliers, a major advantage of truncation to 1-bit coefficients. An application-specific design could likely benefit more from the adder-less $\lambda = \sqrt{2}/2$ pruned correlations, since it is not limited by the static embedded structure of an FPGA, although the 17% ALM reduction from (2) to (3) is notable. Model (4) provides the most dramatic hardware reduction, where the 70% ALM reduction from (3) to (4) is on par with a correlator of 1/4th size.

TABLE 1. FPGA hardware resource utilization.

Model	ALMs	ALUTs	Registers	DSP Blocks
1	79000	144690	180035	4200[1]
2	72801	145281	117985	0
3	62151	106257	112321	0
4	18698	27280	32301	0

[1] Increased from an initial 1400 DSP blocks operating at more than 3× real-time clock speed.

Also of interest is the measured preamble detection performance for the hardware prototypes. All of the non-correlator modules were synthesized using the same Verilog hardware description language (HDL) source, including the actual phase / frequency offset estimator circuits. The thresholding scheme does behave slightly different between variants - to ensure accurate results, the trigger level was set by

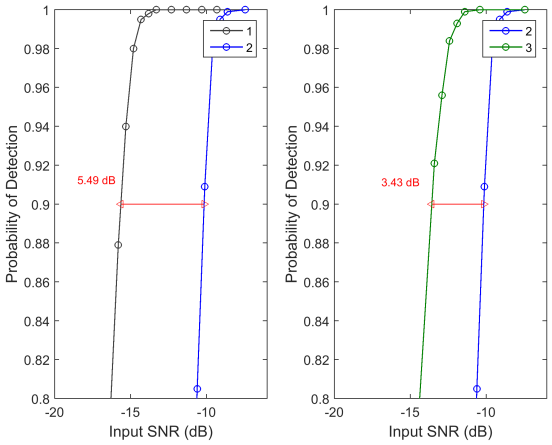

FIGURE 6. Measured implementation loss of truncation to 1-bit coefficients and $\lambda = \sqrt{2}/2$ pruning degradation reduction.

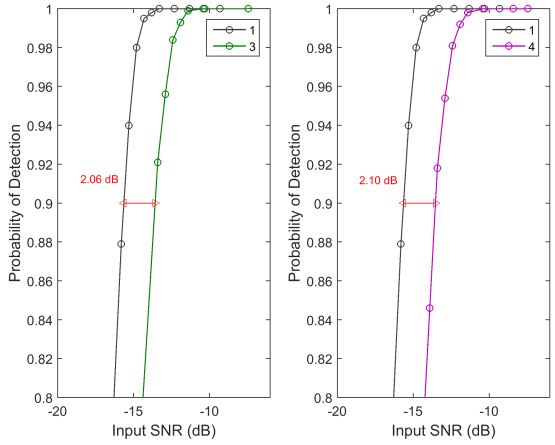

FIGURE 7. Measured implementation loss of the unfolded and folded correlation tap structure with $\lambda = \sqrt{2}/2$ pruned 1-bit truncated coefficients.

empirically searching for a threshold that gives the best performance without returning false detections.

The measured probability of detection (P_D) is shown with respect to signal-to-noise ratio (SNR) in Fig. 5 with the degradation at $P_D = 0.9$ highlighted in Fig. 6 and Fig. 7. As expected, the computationally reduced models do yield reduced performance, yet in a very controlled manner. Given the transformation of complex multipliers to sign-selected adder trees, the loss of 5.49 dB for (2) is tolerable. Model (3) reduces the degradation by 3.43 dB, demonstrating the inherent noise cancellation properties of the amplitude-selective collapse of truncated coefficients. The 2.10 dB performance loss of (4) is the most promising, offering performance almost identical to (3), while providing an overall 76% ALM reduction, requiring 82% fewer dedicated registers, and using no DSP blocks.

V. CONCLUSION

This paper proposed a variety of candidate improvements to the brute-force fallthrough correlator structure, allowing significant computational efficiency improvements and hardware utilization reductions with minimal degradation to preamble detection performance. The truncation of multi-bit precision matched-filter coefficients to 1 bit offers a consolidation of FPGA resources from the brute-force 4200 embedded multipliers and 79000 ALMs to no multipliers and 72801 ALMs. The amplitude-selective collapse of complex-valued coefficients into a single real or imaginary correlation tap further reduces the hardware logic for an overall detection loss of 2.06 dB. Achieving the most substantial hardware reductions is the 4× folded structure with pipelined detection decisions, using no multipliers and 76% fewer ALMs overall, for the trade of only a 2.10 dB performance loss. Moreover, this approach is completely extensible to the Gaussian-shaped digital chaotic spread spectrum signals.

The processing of outputs from computationally reduced correlators needs to consider the expected correlation peak loss in received signal strength estimations, while phase / frequency offset estimations will also be less accurate. By performing an on-time accumulation of the detected preamble signal, any performance loss can be mitigated at the cost of a single shared full-precision multiply-accumulate circuit and some added processing latency. The correlation of shorter preambles will increase estimation error, as necessary for the folded correlation tap structure and largest hardware logic reduction shown in this paper. In that case, storing sub-correlation peak outputs in appropriately sized reference registers is a potential solution, and the optimal sizing of these registers based on the allowable probability of false accept per stage is considered for future work.

ACKNOWLEDGMENT

Thank you to Scott Cowling and Charlie Mesarosh of Zephyr Engineering for implementing the initial truncated coefficients Verilog prototype.

REFERENCES

[1] J. Lei and T.-S. Ng, "Pilot-tone-based maximum likelihood estimator for carrier frequency offset in OFDM systems," *Proc. IEEE Int. Conf. Commun.*, May 2003, pp. 2046–2050.

[2] D. Yan and P. Ho, "Code acquisition in a CDMA system based on Barker sequence and differential detection," in *Proc. 6th Int. Symp. Pers., Indoor Mobile Radio Commun.*, Toronto, ON, Canada, Sep. 1995, pp. 233–236.

[3] C. Du, H. Zeng, W. Lou, and Y. T. Hou, "On cyclostationary analysis of WiFi signals for direction estimation," in *Proc. IEEE Int. Conf. Commun. (ICC)*, Jun. 2015, pp. 3557–3561.

[4] G. E. Corazza and A. Vanelli-Coralli, "Burst vs. Continuous pilot acquisition in wideband CDMA cellular mobile systems," in *Proc. IEEE Wireless Commun. Netw. Conf.*, vol. 3, Sep. 1999, pp. 1080–1084.

[5] S. J. Lee and J. Ahn, "Acquisition performance improvement by Barker sequence repetition in a preamble for DS-CDMA systems with symbol-length spreading codes," *IEEE Trans. Veh. Technol.*, vol. 52, no. 1, pp. 127–131, Jan. 2003.

[6] W. S. Yuan and C. N. Georghiades, "Rapid carrier acquisition from baud-rate samples," *IEEE Trans. Commun.*, vol. 47, no. 4, pp. 631–641, Apr. 1999.

[7] J. Lindenlaub and K. Chen, "Performance of matched filter receivers in non-Gaussian noise environments," *IEEE Trans. Commun. Technol.*, vol. 13, no. 4, pp. 545–547, Dec. 1965.
[8] M. K. Sust and A. Goiser, "A combinatorial model for the analysis of digital matched filter receivers for direct sequence signals," in *Proc. IEEE Global Telecommun. Conf. Exhib. 'Commun. Technol.*, vol. 3, Nov. 1989, pp. 1634–1640.
[9] G. Heidari-Bateni and C. D. McGillem, "A chaotic direct-sequence spread-spectrum communication system," *IEEE Trans. Commun.*, vol. 42, no. 234, pp. 1524–1527, Feb. 1994.
[10] A. J. Michaels and D. B. Chester, "Efficient and flexible chaotic communication waveform family," in *Proc. MILCOM*, Nov. 2010, pp. 1250–1255.
[11] A. Martin, Y. Hasan, and R. M. Buehrer, "Physical layer security of hybrid spread spectrum systems," in *Proc. IEEE Radio Wireless Symp.*, Austin, TX, USA, Jan. 2013, pp. 370–372.
[12] A. J. Michaels and D. B. Chester, "Adaptive correlation techniques for spread spectrum communication systems," in *Proc. IEEE Mil. Commun. Conf.*, Baltimore, MD, USA, Nov. 2016, pp. 678–681.
[13] A. A. Zaher, "An improved chaotic shift keying technique," in *Proc. 5th Int. Symp. Commun., Control Signal Process.*, May 2012, pp. 1–4.
[14] R. Chakravarthy, K. Huang, L. Zhang, and Z. Wu, "Primary user authentication of cognitive radio network using underlay waveform," in *Proc. Cogn. Commun. Aerosp. Appl. Workshop (CCAA)*, Cleveland, OH, USA, Jun. 2017, pp. 1–5.
[15] E. Petrosky and A. Michaels, "Network scalability comparison of IEEE 802.15.4 and receiver-assigned CDMA," *IEEE Internet Things J.*, vol. 6, no. 4, pp. 6060–6069, Aug. 2019.
[16] A. J. Michaels, "High-order PSK signaling (HOPS) techniques for low-power spread spectrum communications," in *Proc. IEEE 19th Int. Symp. World Wireless, Mobile Multimedia Netw.*, Jun. 2018, pp. 01–07.
[17] M. B. Pursley, T. C. Royster, and M. Y. Tan, "High-rate direct-sequence spread spectrum," in *Proc. IEEE Mil. Commun. Conf.*, vol. 2, Feb. 2003, pp. 1101–1106.
[18] R. M. Hewlitt and E. S. Swartzlantler, "Canonical signed digit representation for FIR digital filters," in *Proc. IEEE Workshop SIGNAL Process.*, May 2000, pp. 416–426.
[19] F. J. Harris, C. Dick, and M. Rice, "Digital receivers and transmitters using polyphase filter banks for wireless communications," *IEEE Trans. Microw. Theory Techn.*, vol. 51, no. 4, pp. 1395–1412, Apr. 2003.
[20] A. Michaels, "Digital chaotic communications," Ph.D. dissertation, School Elect. Comput. Eng., Georgia Inst. Technol., Atlanta, Georgia, 2009.