Zero-Order Reconstruction Optimization of Waveforms (ZOROW) for Modest DAC Rates

Charles A. Mohr[1,2], Jonathan W. Owen[1], Shannon D. Blunt[1], Christopher T. Allen[1]
[1] Radar Systems Lab (RSL), University of Kansas, Lawrence, KS
[2] Sensors Directorate, Air Force Research Laboratory (AFRL), Wright-Patterson AFB, OH

***Abstract*—The use of waveforms possessing constant envelope and good spectral containment, combined with high-fidelity arbitrary waveform generation (AWG) capability, has been shown to minimize transmitter distortion effects when waveform diversity is needed. Software-defined radios (SDRs) represent a lower cost, attritable alternative to AWGs at the price of lower fidelity. Specifically, the lower digital-to-analog conversion (DAC) rates supported by SDRs introduce new design challenges for the generation of waveforms requiring high fidelity, most notably those containing in-band spectral notches. Here these challenges are addressed via the ZOROW waveform design scheme that accounts for physical attributes of the SDR. Distortion effects arising from the modest DAC rate are characterized and their mitigation is demonstrated using experimental measurements.**

Keywords—waveform diversity, FM noise, pulse agility, software-defined radar

I. Introduction

A given radar waveform can only be as good as the system used to produce it. The first step beyond the idealistic representation of simulation tends to be high-performance test equipment in the laboratory, where AWGs having DAC rates on the order of GHz are increasingly common. These devices offer tremendous flexibility in terms of spectral shaping and waveform generation that make DAC distortion effects essentially unnoticeable, in some cases even removing the need for up-conversion altogether by permitting waveforms to be produced directly at passband. The key to these capabilities compared to a more modest system is a DAC rate that is many times greater than the waveform's 3-dB bandwidth, a higher effective number of bits (ENOB), a greater analog bandwidth, etc., therby resulting in an excellent approximation of the intended signal.

Deployment in large quantities necessitates a different perspective, however. In recent years there has been increasing interest in the use of SDRs as software defined radars (also SDR) for such purposes as cognitive radar [1-3], automotive radar [4], medical imaging [5], dual-polarized radar [6], synthetic aperture radar (SAR) [7], multiple-input multiple output (MIMO) radar [8, 9], and more. Continued advances in digital and radio frequency (RF) technology have made SDRs of this type more capable and less expensive, thus making them attractive platforms from which to evaluate new radar techniques in a variety of environments. Of course, the lower fidelity relative to AWGs bears consideration, particularly when high fidelity is required.

Here we wish to facilitate the generation of waveforms containing deep spectral notches, which necessitates a high fidelity implementation, using a modest fidelity SDR platform. Such waveforms have themselves been the subject of a growing body of research as a prospective means with which to enable spectrum sharing in a cognitive manner [10, 11]. The fidelity dichotomy is addressed by properly incorporating the physical system attributes of the SDR into the waveform design/generation process denoted as ZOROW. A companion paper [2] leverages this process to implement a real-time cognitive radar spectrum sharing capability on an SDR.

II. DAC Operation and Waveform Model

As the name implies, the purpose of a DAC is to convert a stream of digital values into a continuous, analog signal. In the primary step of this process, the DAC outputs a sequence of voltages that are proportional to the input digital values. More specifically, the input samples to a DAC can be modeled as a train of impulses separated by $T_s = 1/f_s$, where f_s is the DAC rate. While there are other implementations tailored for direct implementation of signals at higher Nyquist zones [12], the DAC typically produces a voltage proportional to the input sample and holds this value for T_s seconds before repeating the process for the next sample. In this way, the DAC reconstructs a continuous time signal in a *zeroth-order hold* manner.

This reconstruction process can be modeled as a sequence of contiguous, rectangular signal structures, where the n^{th} rect(·) function is scaled by the n^{th} value in the digital sequence, denoted as d_n. For convenience and for application to pulsed radar, let the resulting signal be time-limited such that

$$s(t) = \begin{cases} \sum_{n=1}^{N} d_n \,\mathrm{rect}\left(\frac{t - T_s(n - 1/2)}{T}\right) & 0 \le t \le T \\ 0 & \text{otherwise} \end{cases}, \tag{1}$$

where

$$\mathrm{rect}(t) = \begin{cases} 1 & -1/2 \le t \le 1/2 \\ 0 & \text{otherwise} \end{cases}. \tag{2}$$

The construction via (1) produces a signal that is $T = NT_s$ seconds long. Additionally, for an RF system implementing signals at complex baseband, two reconstructions occur simultaneously for the in-phase (I) and quadrature-phase (Q) components. Both I and Q are conveniently represented by (1)

This work was supported in part by the Office of Naval Research under Contract #N00014-16-C-2029 and by the Army Research Office under Contract #W911NF-15-2-0063. DISTRIBUTION STATEMENT A. Approved for Public Release.

by assuming d_n is a complex number having some magnitude and phase.

The simple form of (1) is readily amenable to determining the analytical spectrum via Fourier transform, yielding

$$S(f) = \frac{\sin\left(\pi f T_{\mathrm{s}}\right)}{\pi f} \sum_{n=1}^{N} d_n \exp\left(-j2\pi f (n - 1/2) T_{\mathrm{s}}\right). \quad (3)$$

Thus the baseband spectrum of a DAC-implemented signal is a linear combination of complex exponentials, each weighted by a complex sample in the digital sequence, such that each spectral interval $[(m-1/2)f_{\mathrm{s}}, (m+1/2)f_{\mathrm{s}}]$ for $m \in \mathbb{Z}$ is a copy (or "image") of the fundamental interval $[-f_{\mathrm{s}}/2, +f_{\mathrm{s}}/2]$. This spectral content pertaining to the summation portion of (3) is subsequently scaled by the sinc(·) envelope in the front of (3). Consequently, the power spectrum is only guaranteed to roll-off as steeply as a $\text{sinc}^2(\cdot)$ function. Modern DACs address this otherwise poor spectral containment in a variety of ways that incur some trade-offs.

A. DAC spectral manipulation

Before constructing the analog signal in (1), some DACs employ digital filters to manipulate the digital signal's spectrum, with the following serving as a brief review. For example, the signal can be shaped via a $\text{sinc}^{-1}(\cdot)$ filter, which as the name implies, pre-distorts the digital spectrum such that the sinc(·) envelope scaling in (3) realizes only the summation portion of the baseband spectrum.

Digital interpolation filters can also interpolate a digital sequence at the lower DAC input rate $\bar{f}_s$ by some factor such that it matches f_{s}, the output DAC rate in (1). Critically, this interpolation is performed such that the interpolated sequence possesses approximately the same bandwidth (3-dB for instance) as the original sequence. If sufficient bandwidth is available, the DAC can additionally pre-shift the signal from baseband to some fraction of the DAC rate to assist in the removal of local oscillator (LO) leakage after up-conversion since the resulting analog signal would then be offset from the LO frequency.

Overall, interpolation has the effect of expanding the frequency axis so that more room is created between the baseband spectrum and its images. Consequently, it becomes easier for the analog reconstruction (or "image rejection") filter, which is applied after the signal in (1) is formed, to isolate the desired baseband spectrum prior to up-conversion. These procedures are generally well-known, but it is important to consider their impact on the generation of high-fidelity radar waveforms. More information on the analog implementation of digital signals can be found in [13].

B. DACs and radar signals

Two properties that often differentiate radar signals from others is high transmit power and wide bandwidth. To achieve the highest power level and efficiency, radar signals are generally designed to have a constant envelope such that they can pass through amplifiers operating in saturation with minimal distortion. To maximize the power output from a DAC, this condition also means utilizing the full-scale output (maximum amplitude) of the DAC.

As far as spectral content is concerned, it is not unusual for radar signals to employs hundreds of MHz, or even GHz, of bandwidth. For a modest DAC rate system, such as a commercial SDR where $\bar{f}_s$ is on the order of 100 MHz or so, this need incentivizes the use of as much of $\bar{f}_s$ for signal content bandwidth as possible; a condition that can have adverse effects on achievable fidelity.

C. DAC distortion effects

First consider the analytical spectrum in (3). Due to the shape of the $\text{sinc}^2(\cdot)$ envelope in the associated power spectrum, the signal power contained in the images is disproportionately higher near $f_{\mathrm{s}}/2$. In an extreme case, where the spectrum of the digital sequence is perfectly flat (and normalized for convenience), it is found that

$$100\% \times \int_{-1/2}^{1/2} \left(1 - \frac{\sin^2(\pi f)}{(\pi f)^2}\right) \approx 22.63\% \quad (4)$$

of the signal power resides entirely in the image components. Even more problematic, application of a $\text{sinc}^{-1}(\cdot)$ filter actually increase the signal power content around $\pm f_{\mathrm{s}}/2$, thus leading to even more signal power residing in the images.

The analog reconstruction filter is then tasked with dissipating this image power, which leads to heating (and subsequent dissipation requirements). Moreover, the lower output power produced by the DAC leads to lower overall power efficiency. The digital interpolation filters can help to address these problems, though they introduce their own issues.

Now rewrite the signal in (1) based on a constant amplitude digital sequence at an input DAC rate of $\bar{f}_s$, such that

$$s(t) = \begin{cases} \sum_{n=1}^{N} \exp(j\phi_n)\, \text{rect}\left(\frac{t - \bar{T}_{\mathrm{s}}(n-1/2)}{T}\right) & 0 \le t \le T \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

where $\bar{T}_{\mathrm{s}} = 1/\bar{f}_s$, and each d_n is now unit amplitude and fully described by its phase ϕ_n. Therefore the analytical spectrum in (3) becomes

$$S(f) = \frac{\sin\left(\pi f \bar{T}_{\mathrm{s}}\right)}{\pi f} \sum_{n=1}^{N} \exp\left(j(\phi_n - 2\pi f (n-1/2)\bar{T}_{\mathrm{s}})\right). \quad (6)$$

When this signal is passed through the $\text{sinc}^{-1}(\cdot)$ filter and a digital interpolation filter that increases $\bar{f}_s$ to f_s, it invariably realizes some degree of amplitude modulation (AM). If the signal was already set to the DACs full-scale output, this AM effect introduces signal components that exceed the full-scale value, ultimately resulting in distortion due to clipping. This problem can be mitigated by reducing the input signal amplitude from the full-scale output of the DAC, but doing so likewise reduces the output power from the DAC.

To address these assorted issues, the Zero-Order Reconstruction Optimization of Waveforms (ZOROW) approach was developed. This implementation-oriented design scheme is intended for use with modest DAC-rate systems.

III. Zero-Order Reconstruction of Waveforms

Since, the ZOROW approach is based on the signal structure of (5), ZOROW signals are designed at the input DAC rate. The purpose of ZOROW is to compensate, to the degree possible, for the distortion effects discussed in the previous section.

Given the signal model of (5), define the phase sequence

$$\boldsymbol{\phi} = [\,\phi_1 \;\; \phi_2 \;\; \cdots \;\; \phi_N\,]^T \tag{7}$$

corresponding to the discretized sequence

$$\mathbf{s} = \exp(j\boldsymbol{\phi}) \tag{8}$$

that we wish to convert into an analog signal with minimal distortion. It is important to note that the spectral content of the resulting analog signal (even under idealistic conditions) is not simply the discrete Fourier transform (DFT) of (8), but is instead the analytical (and continuous) spectrum from (6).

In [14], the means to manipulate the analytical spectrum of a first-order representation was proposed by exploiting the fact that sampling theory is reciprocal. In other words, the Nyquist (perfect) reconstruction principle likewise applies when discretizing in the frequency domain as long as the "frequency sampling rate" (the inverse of the separation between frequency samples) is at least twice the temporal duration of the signal, which is a finite value in the case of a pulsed radar waveform. Under this condition, the continuous spectrum in (6) can be perfectly reconstructed from

$$S(f_m, \boldsymbol{\phi}) = \frac{\sin\left(\pi f_m \overline{T}_{\mathrm{s}}\right)}{\pi f_m} \sum_{n=1}^{N} \exp\left(-j(2\pi f_m (n-1/2)\overline{T}_{\mathrm{s}} + \phi_n)\right), \tag{9}$$

where

$$f_m = m\Delta f \tag{10}$$

for integer m on the interval $-\infty < m < \infty$ so long as

$$\Delta f \le \frac{1}{2T}. \tag{11}$$

Since index m has an infinite number of values, representation of (9) to facilitate manipulation of the signal necessitates truncation. Conveniently, all frequency intervals beyond $-\overline{f}_{\mathrm{s}}/2 \le f \le +\overline{f}_{\mathrm{s}}/2$ in either (6) or (9) are only images of the fundamental interval to within a $\mathrm{sinc}(\cdot)$ envelope scale factor. Thus the entire spectrum can be unambiguously represented by discretizing this fundamental interval, for which the minimum number of samples is

$$M = 2T\,\overline{f}_{\mathrm{s}} - 1, \tag{12}$$

where T must be chosen such that $T\overline{f}_{\mathrm{s}}$ (the "time/DAC-rate product") is an integer. If the baseband bandwidth is on the order of the DAC rate, then this term is also relatable to the well-known time/bandwidth product (BT) of a waveform.

IV. ZOROW Spectral Notching

Generally speaking, the type of radar waveforms that will be the most adversely affected by the DAC-induced distortion effects discussed above are those necessitating the highest degree of precision; namely, those that could achieve extremely low range sidelobes (e.g. [15, 16]) and those realizing deep spectral notches [17]. Here we consider the latter application.

To address this fidelity limitation for spectral notching, we shall use the cost function

$$J = \sum_m |\,S(f_m;\boldsymbol{\phi})\,|^2 \tag{13}$$

specified in [14], where f_m for particular values of m span some interval(s) $[f_{\min}, f_{\max}]$ where spectral notch(es) are required, according to the frequency spacing in (11). Therefore (13) represents the summation of signal power in these intervals, with the goal being the determination of $\boldsymbol{\phi}$ such that (13) is minimized. Note that this cost function has relatively little impact on the signal spectrum outside of the formation of spectral notches, and thus should be applied after other waveform design measures have been employed.

Since this cost function is nonlinear and non-convex, gradient descent methods are used to locally minimize (13), relying on the "good enough optimality" of random FM waveforms [18]. These iterative methods take advantage of the current cost function value and gradient, along with previous search directions, to adjust $\boldsymbol{\phi}$ such that the cost function decreases at each iteration. The update is performed as

$$\boldsymbol{\phi}_{k+1} = \boldsymbol{\phi}_k + \mu_k\,\mathbf{p}_k \tag{14}$$

where μ_k is the step size based on a simple backtracking technique [19], and $\mathbf{p}_k$ is the search direction at the k^{th} iteration according to

$$\mathbf{p}_k = \begin{cases} -\boldsymbol{g}_0 & \text{when } k = 0 \\ -\boldsymbol{g}_k + \beta\mathbf{p}_{k-1} & \text{otherwise} \end{cases}, \tag{15}$$

with $0 < \beta < 1$ based on the heavy ball gradient [20].

The gradient of (13) with respect to $\boldsymbol{\phi}$ is calculated to be

$$\nabla_{\boldsymbol{\phi}} J = 2\Im\left\{\sum_m \left(\nabla_{\boldsymbol{\phi}} S(f_m;\boldsymbol{\phi})\right)^* S(f_m;\boldsymbol{\phi})\right\}, \tag{16}$$

where $\Im\{\cdot\}$ extracts the imaginary part of the argument and

$$\nabla_{\boldsymbol{\phi}} S(f_m;\boldsymbol{\phi}) = \begin{bmatrix} \dfrac{\partial S(f_m;\boldsymbol{\phi})}{\partial \phi_1} & \cdots & \dfrac{\partial S(f_m;\boldsymbol{\phi})}{\partial \phi_N} \end{bmatrix}^T. \tag{17}$$

Finally, for a single value ϕ_n, the partial derivative of f_m is

$$\frac{\partial S(f_m;\boldsymbol{\phi})}{\partial \phi_n} = j\,\frac{\sin(\pi f_m \overline{T}_{\mathrm{s}})}{\pi f_m} \exp\left(j(\phi_n - 2\pi f_m (n-1/2)\overline{T}_{\mathrm{s}}\right), \tag{18}$$

where the $\mathrm{sinc}(\cdot)$ spectral envelope shaping by the DAC is naturally included.

V. Experimental Results

We now assess how ZOROW and other physically realizable implementations of notched radar waveforms perform. These methods are evaluated first in simulation, then on a high-fidelity AWG, and finally on an Ettus X310 SDR.

A. Test cases

We consider the generation of spectrally-notched random FM waveforms for the application in [17]. Two sets of 100 unique FM waveforms are produced, the first set having a single notch that is 10% of the 3-dB bandwidth and the second set having an additional notch that is symmetric about the center frequency. The dual notch case reveals the contribution of IQ imbalance on notch depth as explained in Table 1.

The two waveform sets are initialized using the pseudo-random optimized (PRO) FM approach of [21], for which relatively shallow spectral notches can be achieved [17]. Each of these waveforms is then modified using ZOROW to provide greater notch depth compatible with an SDR platform.

For comparison, the set of 100 single-notch PRO-FM waveforms were also modified using the analytical spectrum notching (ASpeN) approach [14] that relies on the polyphase-coded FM (PCFM) implementation [22], which is a *first-order hold* representation when using a rectangular shaping filter. The ASpeN scheme has been experimentally demonstrated to achieve spectral notches with depths exceeding 55 dB when using a high-fidelity AWG.

Additional practical effects being assessed (beyond IQ imbalance and modest DAC rate) include spectral containment, pulse rise/fall-time, and full-scale voltage clipping. Spectral containment is addressed by placing another pair of 10% notches at each band edge. To assess the impact of pulse rise/fall-time effects the inclusion (or not) of a Tukey taper on the pulse shape is considered. Finally, aside from one case to demonstrate the degradation incurred, all waveforms were implemented at 50% of the full-scale DAC voltage. Table I delineates the test cases being considered. The bold/underlined portion is used to name the waveform set in later results.

TABLE I: TEST CASES

Test Case	Demonstration Goal
<u>ASpeN</u> - Single notch, Tukey tapered pulse	ASpeN produces deep spectral notches on high fidelity hardware. Reveals loss in notch depth due to model mismatch at modest DAC rate.
ZOROW - **<u>Single notch</u>**, Tukey tapered pulse	ZOROW waveforms are designed specifically for modest DAC rates. Demonstrates excellent notch depth on the SDR.
ZOROW - **<u>Dual notch</u>**, Tukey tapered pulse	At baseband, IQ imbalance fills in notches with power from the opposite side of zero frequency. Symmetric notches remove this distortion, illustrating the potential given IQ compensation.
ZOROW - Single notch, **<u>Rect pulse</u>** (no taper)	The fast rise and fall times of constant modulus waveforms lead to signal distortion on the SDR if not tapered. Taper removed to show this effect.
ZOROW - Single notch, Tukey tapered pulse, **<u>full-scale</u>** output	Filtering leads to clipping (nonlinear distortion) when the full scale DAC voltage is exceeded.

B. Simulated baseline

The root-mean-square (RMS) spectra is calculated over the 100 single-notch waveform sets for ASpeN and ZOROW using their respective analytical spectrum models, the latter being (6). Figure 1 illustrates these spectra along with the RMS spectra computed over the 100 PRO-FM single-notch waveforms. Figure 2 likewise shows the RMS autocorrelations for these three waveform sets.

Figs. 1 and 2 can be considered ideal in the sense that any subsequent distortion will result in diminished performance by comparison. The ratio between the number of samples in each waveform (1000 here) and BT places the 3-dB point at approximately ± 0.25 in normalized frequency. Interestingly, in simulation the ZOROW waveforms achieve a lower notch depth and retain lower sidelobe levels in comparison to ASpeN. As a final note, the notch edges (see Fig. 1) have also been tapered as a means to reduce the $\sin(x)/x$ range sidelobes that otherwise occur when spectral notches have sharp edges (i.e. rectangular) [17].

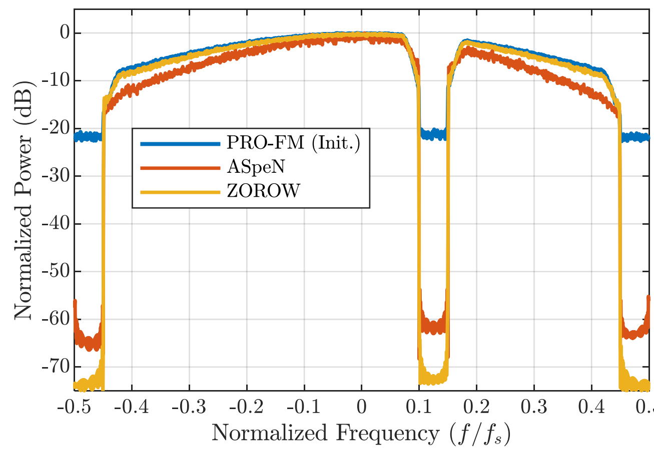

Fig.1. Simulated RMS spectra of the single-notched ASpeN and ZOROW waveforms and their PRO-FM initialization

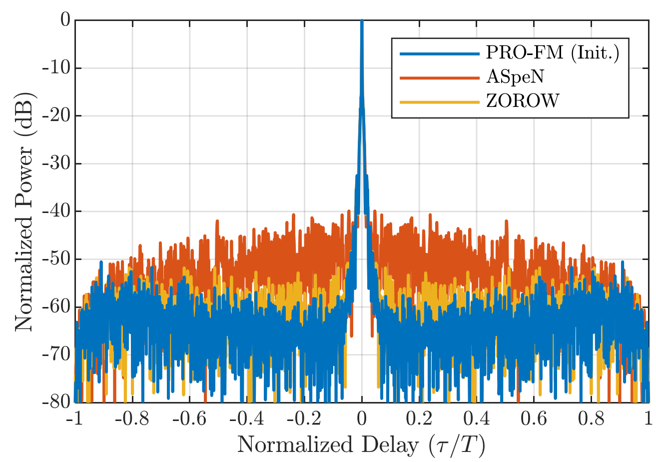

Fig. 2. Simulated RMS autocorrelations of the single-notched ASpeN and ZOROW waveforms and their PRO-FM initialization

C. Experimental results – AWG loopback

The sets of single-notch waveforms were then implemented on a Tektronix AWG at a 3-dB bandwidth of 100 MHz and subsequently captured by a Rohde & Schwarz real-time spectrum analyzer (RSA) connected in loopback at a receive sample rate of 200 MSamples/s. Since the AWG has a high DAC rate (10 GSamples/s), the signals were generated directly at a center frequency of 2 GHz by first performing up-sampling and digital up-conversion using MATLAB™. As shown in Fig.

3, and demonstrated in [14], the ASpeN waveforms are well suited to this implementation and realize a notch depth greater than 50 dB. In contrast, the ZOROW waveforms realize a significant loss in notch depth, though it is still better than 40 dB. Figure 4 shows the autocorrelation results to be virtually identical to Fig. 2.

The reason for notch depth degradation is because, from an AWG perspective, the ZOROW waveform is just a phase code with a chip time of $\overline{T}_s$, which is many time greater than the AWG's DAC period of T_s. Thus, to perform this implementation and preserve the baseband ZOROW spectrum, each ZOROW sequence must be up-sampled in the same manner as a phase code, i.e. each sample was duplicated $L = 50$ times to conform to the DAC rate of the AWG. This baseband signal was then modulated by a 2 GHz carrier frequency. However, when the real part of this passband signal is extracted for implementation on the AWG, the resulting image spectrum produced at −2 GHz exhibited a sin(x)/x spectral roll-off (due to the phase-code structure) that partially filled in the spectral notch, despite being 4 GHz away in absolute bandwidth (i.e. 40 times the 3-dB bandwidth). This effect is not an issue on the SDR (for which ZOROW is intended) because it applies the reconstruction filter prior to up-conversion.

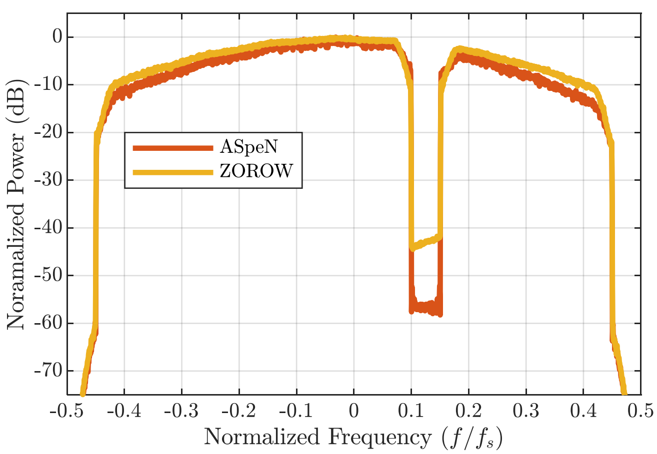

Fig. 3. Experimental RMS autocorrelations of the single-notched ASpeN and ZOROW waveforms implemented on an AWG.

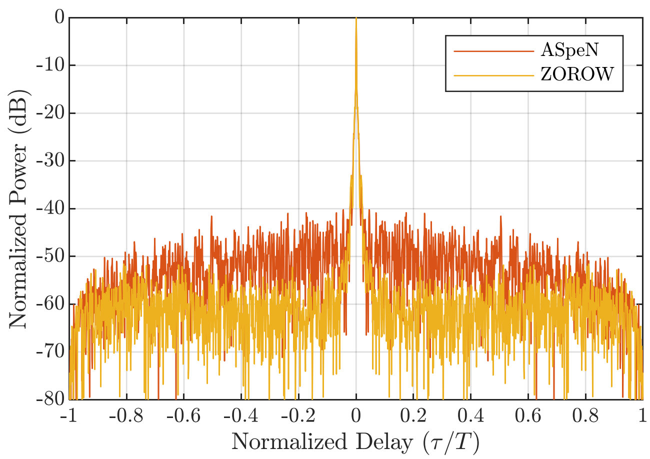

Fig. 4. Experimental RMS spectra of the single-notched ASpeN and ZOROW waveforms implemented on an AWG

D. Experimental results – SDR loopback

The various signal sets were finally implemented on an Ettus X310 SDR [23] at a center frequency of 2 GHz and likewise captured in loopback on the same RSA. The internal X310 DAC employs an internal ×4 digital interpolation filter. Figure 5 depicts the resulting RMS spectra for all cases in Table I, with peak-normalized versions showing a close-up of the single notch illustrated in Fig. 6.

Despite the outstanding results demonstrated on the AWG, the worst performer in terms of notch depth is the ASpeN set of waveforms (green trace). The reason for this seeming discrepancy is because the ASpeN implementation simply cannot be realized with sufficient fidelity on the SDR due to the far lower DAC rate.

The next two sets of waveforms, with almost identical absolute notch levels (though note difference in peak power), are the ZOROW implementations involving no edge tapering of the pulse (yellow trace) and the full-scale DAC version that experiences amplitude clipping (purple trace). While these effects clearly arise from nonlinear distortion, the actual reason for the former is not yet fully known, though it appears to be related to the sharp rise/fall-time.

Finally, the most successful cases are the ZOROW waveforms that possess less than full-scale DAC output and a Tukey taper on the pulse (orange and blue traces). Both cases realize a relative notch depth of more than 50 dB. Moreover, the dual notch case that exploits symmetry (to isolate IQ imbalance) shows roughly 3 dB greater notch depth than the single notch case. While not necessarily a solution to the imbalance problem, this difference does indicate that compensation of this effect could enhance performance.

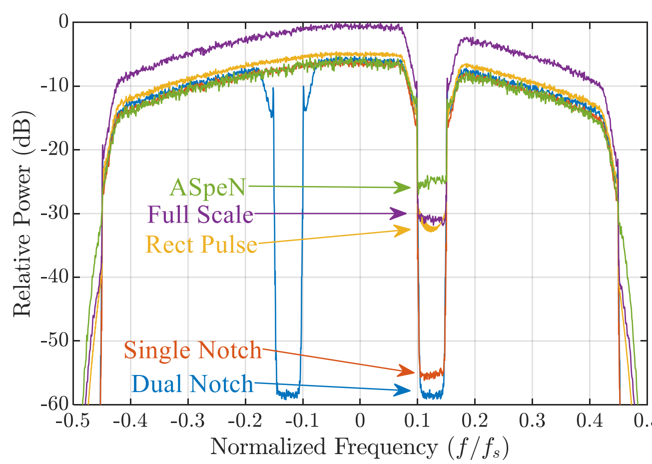

Fig. 5. Experimental RMS spectra of the test cases in Table I implemented on the X310 SDR.

Figure 7 shows the RMS autocorrelations for the cases involving Tukey-tapering of ASpeN, Tukey-tapering of ZOROW (single notch), and ZOROW with no taper. The trade-off incurred for greater notch depth (via the taper) is observed here to be higher range sidelobes, which is not unexpected since tapering of random FM waveforms does not provide the spectral shaping benefit that is observed when tapering a linear FM chirp.

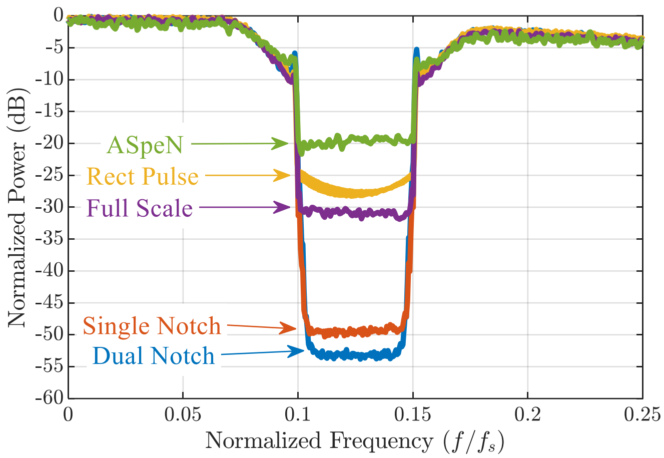

Fig. 6. Peak-normalized experimental RMS spectra of the test cases in Table I implemented on the X310 SDR, with close-up of the notch

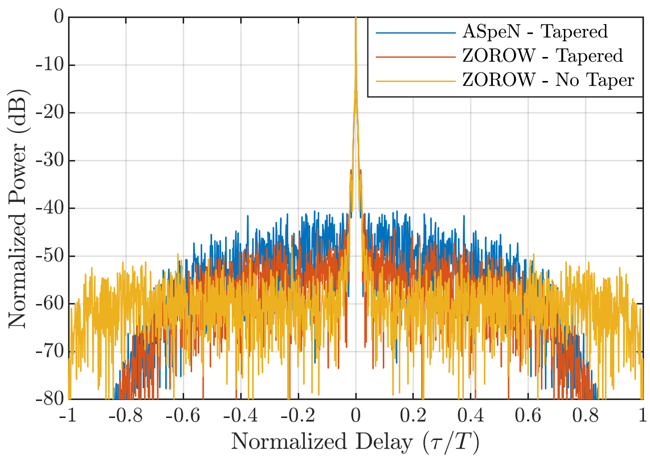

Fig. 7. Experimental RMS autocorrelations of three of the test cases in Table I implemented on the X310 SDR

VI. Conclusions

The growing capabilities and low cost of SDR systems make them attractive platforms for a variety of emerging radar applications. However, the lower fidelity that they presently achieve relative to high-performance AWGs introduces a new set of implementation trade-offs that must be carefully considered. Here the ZOROW waveform modification scheme was proposed and experimentally demonstrated as one prospective way in which this trade-space may be addressed when higher fidelity is needed to achieve transmit spectral notching at the more modest DAC rates of available SDRs. As shown in a companion paper [2], this method likewise facilitates a form of real-time cognitive radar spectrum sharing.

References

[1] B.H. Kirk, R.M. Narayanan, K.A. Gallagher, A.F. Martone, K.D. Sherbondy, "Avoidance of time-varying radio frequency interference with software-defined cognitive radar," *IEEE Trans. Aerospace & Electronic Systems*, vol. 55, no. 3, pp. 1090-1107, June 2019.

[2] J.W. Owen, C.A. Mohr, B.H. Kirk, S.D. Blunt, A.F. Martone, K.D. Sherbondy, "Implementation of spectrally-notched random FM waveforms on a real-time cognitive radar platform," submitted to *IEEE Intl. Radar Conf.*, Washington, DC, Apr./May 2020.

[3] J.M. Christiansen, G.E. Smith, K.E. Olsen, "USRP based cognitive radar testbed," *IEEE Radar Conf.*, Seattle, WA, May 2017.

[4] H. Zhang, L. Li, K. Wu, "24 GHz software-defined radar system for automotive applications," *European Conf. Wireless Technologies*, Munich, Germany, Oct. 2007.

[5] J. Marimuthu, K.S. Bialkowski, A.M. Abbosh, "Software-defined radar for medical imaging," *IEEE Trans. Microwave Theory & Techniques*, vol. 64, no. 2, pp. 643-652, Feb. 2016.

[6] J. Hershberger, T. Pratt, R. Kossler, "Implementations of coherent software-defined dual-polarized radars," *IEEE Trans. Microwave Theory & Techniques*, vol. 65, no. 5, pp. 1673-1681, May 2017.

[7] J. Mar, Y. Lin, "Implementation of SDR digital beamformer for microsatellite SAR," *IEEE Geoscience & Remote Sensing Letters*, vol. 6, no. 1, pp. 92-96, Jan. 2009.

[8] M.T. Frankford, N. Majurec, J.T. Johnson, "Software-defined radar for MIMO and adaptive waveform applications," *IEEE Intl. Radar Conf.*, Washington, DC, May 2010.

[9] S. Baskar, E. Ertin, "A software defined radar platform for waveform adaptive MIMO radar research," IEEE Intl. Radar Conf., Washington, DC, May 2015.

[10] A. Farina, A. De Maio, S. Haykin, *The Impact of Cognition on Radar Technology*, SciTech Publishing, 2017.

[11] S.D. Blunt, E.S. Perrins, *Radar & Communication Spectrum Sharing*, SciTech Publishing, 2018.

[12] M. Guibord, Digital-to-Analog Converter (DAC) Output Response, Texas Instruments, TIPL 4705, 2017.

[13] A Technical Tutorial on Digital Signal Synthesis, Analog Devices, 1999.

[14] C.A. Mohr, S.D. Blunt, 'Analytical spectrum representation for physical waveform optimization requiring extreme fidelity," IEEE Radar Conf., Boston, MA, Apr. 2019.

[15] E. De Witte, H.D. Griffiths, "Improved ultra-low range sidelobe pulse compression waveform design," Electronics Letters, vol. 40, no. 22, pp. 1448-1450, Oct. 2004.

[16] J. Jakabosky, S.D. Blunt, T. Higgins, "Ultra-low sidelobe waveform design via spectral shaping and LINC transmit architecture," IEEE Intl. Radar Conf., Washington, DC, May 2015.

[17] B. Ravenscroft, J.W. Owen, J. Jakabosky, S.D. Blunt, A.F. Martone, K.D. Sherbondy, "Experimental demonstration and analysis of cognitive spectrum sensing & notching," IET Radar, Sonar & Navigation, vol. 12, no. 12, pp. 1466-1475, Dec. 2018.

[18] S.D. Blunt, et al, " Principles & applications of random FM radar waveform design," submitted to IEEE AES Magazine.

[19] J. Nocedal, S. Wright. Numerical Optimization, Springer Science & Business Media, 2006.

[20] E. Ghadimi, R. Feyzmahdavian, M. Johansson, "Global convergence of the heavy-ball method for convex optimization," European Control Conf., Linz, Austria, July 2015

[21] J. Jakabosky, S. D. Blunt, B. Himed, "Spectral-shape optimized FM noise radar for pulse agility," IEEE Radar Conf., Philadelphia, PA, May 2016.

[22] S.D. Blunt, M. Cook, J. Jakabosky, J. de Graaf, E. Perrins, "Polyphase-coded FM (PCFM) radar waveforms, part I: implementation," IEEE Trans. Aerospace & Electronic Systems, vol. 50, no. 3, pp. 2218-2229, July 2014.

[23] Ettus Research, "USRP X300 and X310 X Series," specifications, Jan. 2019