A Pulse Compression Waveform for Improved-Sensitivity Weather Radar Observations

JAMES M. KURDZO

School of Meteorology, and Advanced Radar Research Center, University of Oklahoma, Norman, Oklahoma

BOON LENG CHEONG

Advanced Radar Research Center, University of Oklahoma, Norman, Oklahoma

ROBERT D. PALMER AND GUIFU ZHANG

School of Meteorology, and Advanced Radar Research Center, University of Oklahoma, Norman, Oklahoma

JOHN B. MEIER

Advanced Radar Research Center, University of Oklahoma, Norman, Oklahoma

(Manuscript received 22 January 2013, in final form 9 September 2014)

ABSTRACT

The progression of phased array weather observations, research, and planning over the past decade has led to significant advances in development efforts for future weather radar technologies. However, numerous challenges still remain for large-scale deployment. The eventual goal for phased array weather radar technology includes the use of active arrays, where each element would have its own transmit/receive module. This would lead to significant advantages; however, such a design must be capable of utilizing low-power, solid-state transmitters at each element in order to keep costs down. To provide acceptable sensitivity, as well as the range resolution needed for weather observations, pulse compression strategies are required. Pulse compression has been used for decades in military applications, but it has yet to be applied on a broad scale to weather radar, partly because of concerns regarding sensitivity loss caused by pulse windowing. A robust optimization technique for pulse compression waveforms with minimalistic windowing using a genetic algorithm is presented. A continuous nonlinear frequency-modulated waveform that takes into account transmitter distortion is shown, both in theory and in practical use scenarios. Measured pulses and weather observations from the Advanced Radar Research Center's dual-polarized PX-1000 transportable radar, which utilizes dual 100-W solid-state transmitters, are presented. Both stratiform and convective scenarios, as well as dual-polarization observations, are shown, demonstrating significant improvement in sensitivity over previous pulse compression methods.

1. Introduction

The current generation of Weather Surveillance Radar-1988 Doppler (WSR-88D) weather radars in the United States is more than 20 years of age (Yussouf and Stensrud 2008). Despite recent major improvements to the network (such as dual-polarization capabilities; Doviak et al. 2000), there are a number of potential enhancements that are currently being explored as researchers look toward the future of weather radar observations. Of key importance to National Weather Service forecasters is the desire for higher temporal resolution in a future weather radar network in order to provide more timely information in rapidly changing weather conditions (Zrnić et al. 2007). The advent of phased array radar (PAR) technology in the weather community has sparked numerous studies regarding the feasibility for a future PAR network in the United States (Weber et al. 2007; Heinselman et al. 2008; Bluestein et al. 2010; McLaughlin et al. 2009; Weadon et al. 2009; Brown and Wood 2012).

Corresponding author address: James M. Kurdzo, Advanced Radar Research Center, 3190 Monitor Avenue, Norman, OK 73019.
E-mail: kurdzo@ou.edu

DOI: 10.1175/JTECH-D-13-00021.1

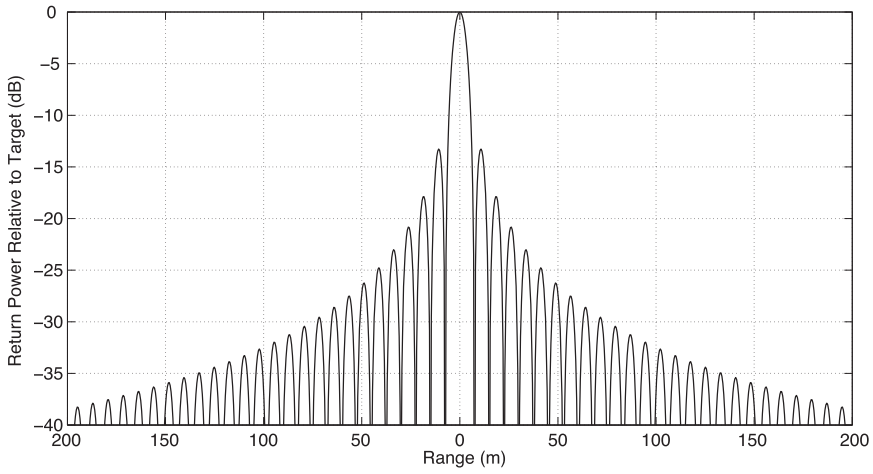

FIG. 1. Compressed pulse ACF of an LFM waveform. This simulates a received pulse being match filtered with a point target. Note the extremely high sidelobes, resulting in a stretching of the target in range.

PAR offers not only rapid volumetric updates, because of the use of electronic steering, but also a number of additional benefits. Active phased array systems use low-power, solid-state transmit/receive elements, resulting in lower overall cost. In fact, low-power transmitters are not a new idea in the weather radar community. The use of solid-state radar transmitters is becoming a popular alternative to high-power klystron, magnetron, and traveling wave tube (TWT) transmitters because of their relatively low cost and improved reliability. The potential benefits of inexpensive, low-power transmitters are wide and varied; however, there are drawbacks associated with their use. Low sensitivity, the principal issue with utilizing low peak power, can be mitigated using a longer pulse with some variant of frequency modulation (i.e., pulse compression).

Pulse compression is a technique that utilizes a modulated transmit waveform in order to decouple the dependency of range resolution from pulse length. Instead of range resolution being directly related to pulse length, it becomes inversely proportional to the bandwidth of the transmit waveform. For frequency modulation, this bandwidth is determined by the range of frequencies covered in the ''chirp'' of the transmit signal. Because of the dependence of range resolution on chirp bandwidth, the pulse length can be increased beyond what may have previously been considered acceptable, without any consequences regarding poor range resolution. For low-power transmitters, such as those discussed above, a long transmit pulse that utilizes pulse compression can theoretically yield acceptable sensitivity values while maintaining acceptable range resolution.

Pulse compression allows for the use of significantly lower-power transmitters while achieving similar performance. Unfortunately, side effects are introduced when using a pulse-compressed waveform. Range sidelobes, a result of autocorrelation when the received pulse is match filtered with the transmit pulse, become detrimental to radar system performance. A linear frequency modulation (LFM) results in peak sidelobes of −13.3 dB, effectively elongating the target in range, causing an artificial target to appear to the user. Figure 1 shows the theoretical autocorrelation function (ACF) of an LFM waveform. Note the high range sidelobes, causing artificial returns in range from a point target.

To combat this problem, target-tracking applications have long used an amplitude-tapered transmit pulse (or ''window'') in order to identify individual targets. Windowing is known to significantly decrease range sidelobe intensity, alleviating the dilemma caused by using pulse compression. However, windowing prevents the transmitter from sending a full-power pulse, since the beginning and end of the pulse are cut significantly in power in order to create the windowed shape. Since many detection-only applications do not require a high level of sensitivity, this method has served as a viable option for inexpensive radar systems using pulse compression. Figure 2a shows the windowing effect on a linear frequency-modulated pulse [windowed linear frequency modulation (WLFM)], while Fig. 2b shows the ACF of the waveform. While the range sidelobes are significantly lower than those in Fig. 1, the integration of power across the pulse shown in Fig. 2a is much lower than the rectangular pulse used to generate the autocorrelation in Fig. 1, causing a loss in sensitivity. In this aggressively windowed case, the loss in sensitivity is approximately 4.6 dB. This loss can be quantified via

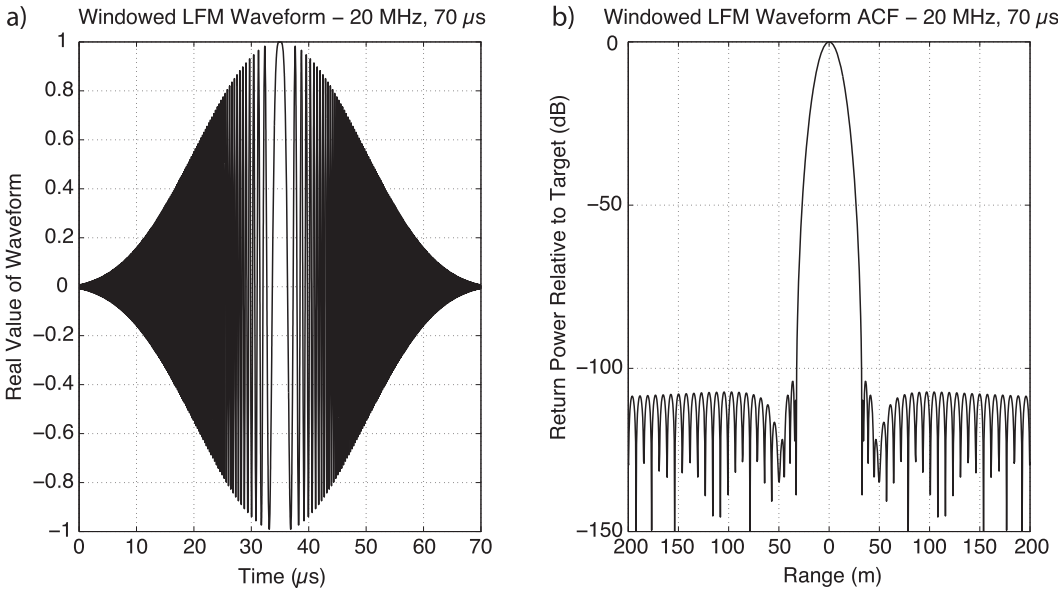

FIG. 2. An example of (a) a heavily windowed transmit pulse and (b) its associated compressed ACF. Note the much lower sidelobes compared with Fig. 1. Because of the heavy windowing, however, significant power loss results.

$$\mathrm{SNR}_{\mathrm{loss}} = 10\log_{10}\left\{\frac{\left(\sum_{n=1}^{N} w_t w_r\right)^2}{N\left[\sum_{n=1}^{N}(w_t w_r)^2\right]}\right\}, \qquad (1)$$

where w_t and w_r represent the amplitude weighting function of the transmit waveform and receive window, respectively; and N is the length of the pulse in samples. The loss is represented in decibels below optimum, which would be achieved with no windowing on transmit or receive. Equation (1) represents a normalized two-way signal-to-noise ratio (SNR) loss that can be formed via a generalization of processing gain methods described in Harris (1978) and Bharadwaj and Chandrasekar (2012) for either matched or mismatched filters.

As interest has grown in using pulse compression for weather radar systems, the method of windowing has come into question. Instead of point targets, weather radar systems deal with distributed targets over a high dynamic range (anywhere from −10 to 75 dBZ). With many windowing methods cutting power efficiency by greater than 3 dB, and as high as 5 dB, weak distributed weather targets suffer greatly from the drop in sensitivity associated with a window, resulting in a lowered SNR and lower overall observability for reflectivity measurements. The loss in SNR leads to the desire to improve sensitivity for a weather radar scenario, meaning a method other than windowing must be used in order to suppress range sidelobes.

This paper describes an optimized frequency-modulated (OFM) pulse compression waveform designed using a genetic algorithm that takes into account individual system characteristics and performance measures in order to design a low SNR loss (high power efficiency) waveform for use with an atmospheric weather radar utilizing pulse compression. In a general sense, the technique presented in this study significantly improves not only theoretical sidelobe performance but also significantly closes the gap between theory and practice in the actual system by tailoring the design to the hardware being used. This method has the potential to be utilized within all types of hardware, ranging from weather radar to military solutions such as airborne- and satellite-based radar systems.

A brief background of pulse compression methods for weather radar is discussed. The framework for the genetic algorithm is presented, followed by the design and testing methodologies. The need for transmitter distortion correction is discussed, and examples are shown with and without predistortion in the optimization algorithm. Results are shown using the theoretical waveform, coupled signals from a real system, and weather observations (in varying scenarios with LFM, WLFM, and OFM waveforms). Finally, a discussion of our findings and a listing of ongoing and future work are presented.

2. Background

While pulse compression has been utilized in military radar applications for many decades, the idea of using pulse compression with weather radar systems is not as well developed and has changed considerably over the past 30 years. The first methods used for weather radar pulse compression appeared in the 1970s and were based

on Barker codes (Ackroyd and Ghani 1973; Austin 1974). While this produced sidelobes lower than a linear frequency chirp, results were quite limited, and the sidelobes were less desirable for the high-sensitivity needs of a weather radar system. After a lull in advancement in the field, a number of new techniques arose in the 1990s and 2000s. Keeler and Hwang (1995) presented one of the first holistic studies involving what was known about pulse compression for weather radar at the time. This was the first mention of the possibility of higher spatial and temporal resolution using phased array antennas for weather observations, and it helped define many of the parameters still used to analyze waveform performance. Also of concern to weather radar specifically was how tolerant pulse compression waveforms would be to Doppler shifts. Bucci and Urkowitz (1993) showed successful simulations indicating a lack of significant degradation using certain types of waveforms.

A number of simulations using Barker codes were performed (Baden and Cohen 1990; Bucci et al. 1997; Mudukutore et al. 1998; Duh et al. 2004), showing limited usefulness for weather radar. Around the same time, Griffiths and Vinagre (1994) presented a study involving the use of a piecewise LFM, also known as nonlinear frequency modulation (NLFM), for use with satellite-based weather radar systems. By introducing slightly different LFM rates at the edges of the waveform, and keeping LFM throughout the center of the waveform, lower sidelobes were made possible. To achieve the desired sidelobes of −62 dB, however, significant windowing was still used in conjunction with the NLFM technique.

De Witte and Griffiths (2004) presented a "continuous NLFM" waveform that improved upon the piecewise NLFM waveform from Griffiths and Vinagre (1994). An amplitude taper was used on receive only, leading to the use of a mismatched filter. While −70-dB sidelobes were achieved in theory, it was not clear how significantly the mismatched filter affected SNR. Further investigation shows the use of a Nuttall window on receive, resulting in expected two-way SNR loss as high as 2.96 dB. Also of significance was their discussion regarding NLFM Doppler tolerance. In cases where the NLFM is symmetric, it was shown that shifts in the ambiguity function with higher Doppler velocities were not as significant as other pulse compression waveform design methods. This was shown again by George et al. (2008), a study that also discussed elements of system calibration and blind range mitigation (which will be covered in later sections of this paper).

Few examples of actual weather data collection using pulse compression exist in the literature. One of the first weather datasets collected using pulse compression was presented in O'Hora and Bech (2007). Despite relatively few examples of weather data collected using pulse compression in the literature, there have been a number of systems developed that utilize pulse compression for weather observations. Weighted NLFM methods have been used recently with low-power, high-frequency radar systems. Carswell et al. (2008) described a solid-state, dual-polarized, dual-wavelength precipitation Doppler radar that operates at Ka band. Additionally, Hirth et al. (2012) presented initial data collected using Ka-band mobile research weather radar systems.

There have also been a limited number of optimization studies regarding pulse compression with weather radar systems, with most focusing on only small details of the waveform, as opposed to a holistic approach that would optimize a waveform from the beginning of the design phase to actual deployment in a system. Among such optimization studies, topics varied from coded waveforms (Ackroyd and Ghani 1973; Baden and Cohen 1990) to window design (Yichun et al. 2005) to mismatch filter design (Cilliers and Smit 2007). Wang et al. (2008) proposed the use of a genetic algorithm to optimize the interpulse frequency step in a stepped LFM; however, they did not explore the possibility of optimizing the entire waveform and minimizing window usage. More recently, Bharadwaj and Chandrasekar (2012) presented a study regarding waveform design for solid-state weather radar systems. While a similar NLFM waveform to De Witte and Griffiths (2004) was used with a mismatched filter, many new elements were presented. Discussions regarding characterization of the SNR loss, as well as the important issue of "blind range" mitigation, offered a refreshed look at the reality of using pulse compression with weather radar.

In the past 8 years, nearly all forms of pulse compression used methods derived from either De Witte and Griffiths (2004) or the stationary phase principle, with various types of windowing (Ge et al. 2008). The best solutions thus far have lacked flexibility to provide low sidelobes without significant windowing. Additionally, while waveforms could be designed with somewhat acceptable power efficiency and theoretically low sidelobes, the implementation through actual systems often results in significantly degraded performance.

3. Methodology

a. *Waveform design and optimization*

The key element to any NLFM waveform is the frequency function. By allowing for a sharper change in frequency along the edges of the frequency function, an effective windowing can be applied to a rectangular waveform. This is because less emphasis is placed on the highest "edge" frequencies. De Witte and Griffiths (2004)

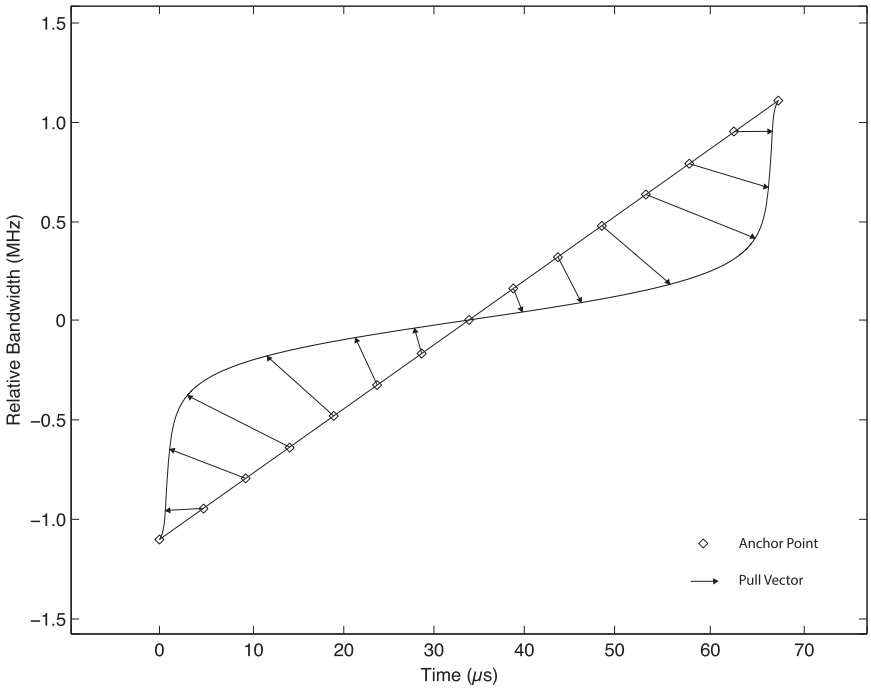

FIG. 3. Demonstration of how a straight line with fixed anchor points can be bent into a complex function using the method of Bézier curves. The anchor points are pulled using a series of vectors. These vectors can be used as variables within an optimization framework in order to design highly flexible frequency functions.

proposed a commonly used method to design an NLFM waveform; however, the flexibility of the function was limited, resulting in a relatively small solution set. The goal for our method is to utilize global optimization techniques in order to match a frequency function to hardware specifications.

While numerous iterations of polynomial-based frequency function representations were attempted, it was found that a more flexible function method was necessary. To provide this flexibility, Bézier curves were employed, which are commonly used in vector graphics applications (Farin 1997). These curves can be implemented in software by defining a straight line and a series of ''anchor points.'' The anchor points can be ''pulled'' using a vector stemming from the line (defined by an X and Y coordinate in a two-dimensional plane). Figure 3 demonstrates the application of Bézier curves to frequency function design. When the anchor points are pulled in various directions, unique shapes can be made from an originally straight line. This significantly increases the overall solution set (and therefore search space within an optimization framework) due to the much higher flexibility for line shapes.

By making use of a Bézier curve for the design of the frequency function, an optimization framework can be developed. For this study, we utilized 15 evenly spaced anchor points along a line spanning the pulse length (along the abscissa) and available bandwidth (along the ordinate). This example is illustrated in Fig. 3. Of the 15 anchor points, the ends and the middle point were locked to the bounds of the design, leaving 12 total changeable anchor points. Because of the desire for symmetry for Doppler tolerance (Bucci and Urkowitz 1993; Griffiths and Vinagre 1994; Levanon and Mozeson 2004; George et al. 2008), the anchor points are mirror images, meaning only six are optimizable. Given that each anchor point contains both an X and Y coordinate, a total of 12 degrees of freedom remain for optimization. Within the optimization, the axes are normalized to the pulse length and available bandwidth, and are sampled at a user-defined resolution that is based on computational availability. For the examples in this paper, a resolution of 2001 points was used in both dimensions, resulting in 2001 possible values for each of the 12 variables. This equates to 4.12×10^{39} possible solutions for the frequency function. It should be noted that depending on the desired performance, as well as the availability of additional computational power, the search space can be much larger. The division of the normalized frequency function into 2001 possible values in each dimension was the feasible resolution for our uses.

Because of the large search space, genetic algorithms were chosen based on their simplicity, flexibility, and speed. Genetic algorithms belong to a class of optimization techniques known as evolutionary algorithms, which are based on the theory of evolution (Eiben and Smith 2007). These types of algorithms work to progressively improve the functionally defined fitness based on evolutionary principles. Genetic algorithms, specifically, are capable of quickly finding global optimum solutions to both simple and complicated nonlinear problems. Genetic algorithms offer one of the most consistent methods to achieve global optimum in a computationally feasible fashion (Holland 1975). While it is possible that other methods could result in slightly more accurate or timely computations, we found that genetic algorithms provided us with successful results in a reasonable amount of time.

Between each generation during the genetic algorithm optimization, a portion of the population with the lowest fitness is discarded, and the remaining population members are randomly paired to create ''children'' (this is termed *crossover*). A key feature of genetic algorithms is the ability to avoid local maxima in fitness. To achieve this, occasional mutations are introduced with a predetermined likelihood of occurrence. The use of mutation, specifically, results in a much higher likelihood that the optimization will not stall at local maxima. Additionally, in order to guarantee the lack of regression in fitness, the top two population members with the highest fitness, or ''elite members,'' are retained through each generation. Generations continue until a maximum fitness score is obtained; this maximum is recognized by a lack of change in the best fitness score for a predetermined number of generations, indicating convergence to a solution, and the end of the optimization.

The genetic algorithm progresses by changing the values of the 12 parameters, which represent the pull vectors of the Bézier anchor points. The predetermined boundary conditions for the optimization are simply the normalized values of pulse length and available bandwidth, which are divided into resolution segments based on the available computational power. Since the actual values of pulse length and bandwidth are not used in the optimization because of normalization, it is the ratio between time and bandwidth that is important in determining the final shape of the frequency function. Each of the 12 parameters contributes to a vectorized ''pull'' of the originally linear frequency function into a nonlinear shape and can be represented in numerical form for operation within the genetic algorithm framework.

While the genetic algorithm uses the 12 degrees of freedom and their predetermined bounds for optimization, there must still be a fitness function, which defines the goal for optimization. Two main factors were determined as being critical in optimization performance: peak sidelobe level (PSL) and main lobe width (MLW). Both indices can be calculated using a theoretical ACF, which simulates the expected matched filter and the associated waveform performance. PSL is defined as the highest point in the ACF outside of the main lobe. It is important to note that this may not always be the first sidelobe. MLW is defined as the null-to-null width of the main lobe and is a proxy both for range resolution and general waveform performance.

The use of MLW within the frequency function, specifically, goes beyond the scope of simply defining waveform performance based on the 3-dB range resolution convention. By focusing on null-to-null MLW, the algorithm performs much more efficiently than using the more common 3-dB MLW. In this manner, the 3-dB range resolution is governed by the bandwidth of the chirp. However, in order to maintain an acceptable main lobe shape, the MLW is constrained within the algorithm. In the examples provided in this paper, a null-to-null MLW of 600 m was set as a constraint so that the main lobe did not become too wide for a weather radar purpose. The fitness function used in the optimization framework is

$$F = \frac{\text{PSL}}{\text{MLW}}, \tag{2}$$

where the units of PSL are in decibels and the units of MLW are in meters.

The genetic algorithm attempts to minimize the fitness function by decreasing PSL and/or decreasing MLW. The algorithm can change these parameters by altering the original 12 degrees of freedom that define the frequency function, then performing an ACF and checking the indices on a successive generation. This process repeats, following standard genetic algorithm procedures, until a stopping criterion is met. Typically, stopping criteria are defined as a minimum change in average fitness, and/or a significant number of generations without an improvement in fitness. Figure 4a offers a visual representation of the genetic algorithm optimization procedure.

A significant advantage to using an optimization technique for waveform design is the ability to build ''pre-distortion'' into the design. Unlike techniques such as the stationary phase principle and other amplitude-modulated waveform design methods, which often produce a single result, optimization allows for an adjustment within the design process, which can account for transmitter imperfections. This is achieved by first transmitting a nonoptimized LFM waveform through the system in question and measuring the coupled pulse. The coupled pulse is recorded in the time series data of each channel for precise matched filtering in processing.

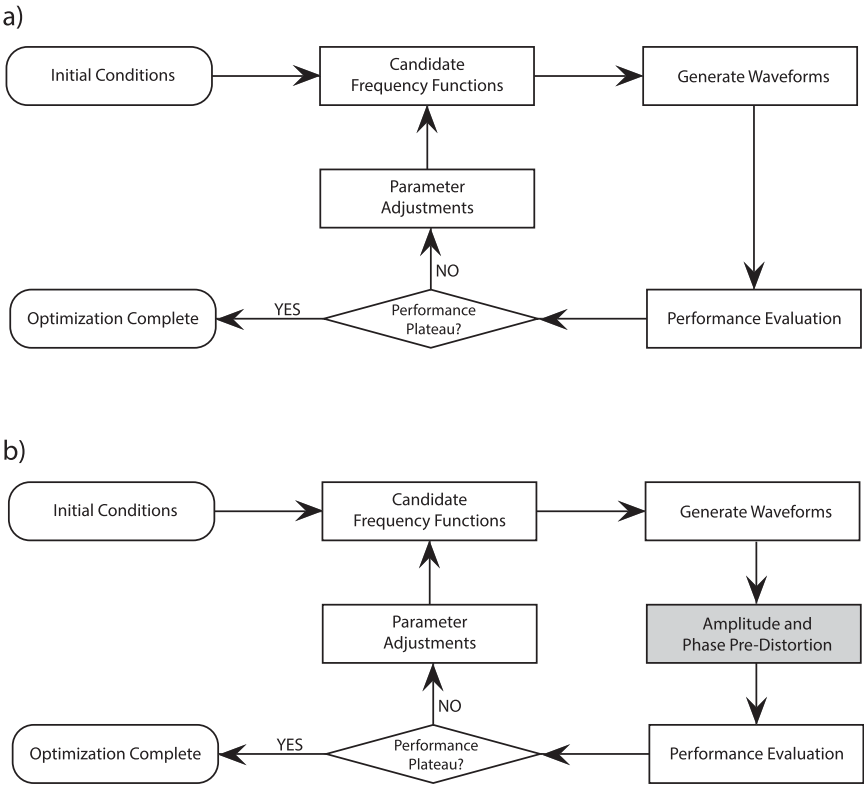

FIG. 4. (a) Flowchart for the genetic algorithm. Potential frequency functions are converted into waveforms based on system specifications, and their theoretical ACFs are tested for performance. If sufficient time has passed without an improvement in performance, based on the specified fitness function, the optimization ends. (b) As in (a), but with amplitude and phase pre-distortion taken into account. The pre-distortion takes measured transmitter and system fluctuations into consideration before testing waveform performance.

The measured pulse, which is affected by transmitter distortions, is then compared with the intended transmit pulse in the frequency domain, allowing for the generation of an effective transfer function of the system. The inverse of the transfer function is then inserted into every calculation of the ACF in the optimization, effectively modeling the expected distortion by the transmitter. Figure 4b shows the added step in the optimization procedure in order to account for the transmitter distortion.

This technique, which has been used in various other system implementations, is an attempt at predicting what will happen to the waveform when it is passed through the nonideal transmitter and can be incorporated into any iterative waveform optimization process. The resulting optimized waveform, when autocorrelated, does not appear to be optimal. However, when it is sent through the actual system, the resulting transmit pulse displays significantly improved performance. This additional method can only be used in an optimization design technique in which the waveform can be dynamically changed from generation to generation. By accounting for expected transmitter distortions, the actual performance in the system becomes much closer to theory.

When combining pre-distortion with the flexibility of the method used in the optimization algorithm, the resulting observed waveform through the real-world system achieves significantly better performance than previous methods. This method can be used with any system capable of using pulse compression and routinely achieves better PSL sidelobe performance than heavily windowed methods while utilizing high power efficiencies with SNR loss as low as 0.05 dB. The resulting sensitivity gain in our system tests is as high as 2.95 dB compared with methods that utilize heavy windowing. Additionally, with the built-in pre-distortion, sidelobes are generally much closer to theoretical designs.

b. *Testing platform*

The platform being used for testing waveform design at the University of Oklahoma is the PX-1000 transportable polarimetric X-band radar system (Cheong et al. 2013). Table 1 describes the main characteristics of the system. PX-1000 operates via two independent

TABLE 1. System characteristics of PX-1000.

Transmitter type	Dual solid-state power amplifiers
Operating frequency	9550 MHz
Sensitivity	$<$20 dBZ at 50 km
Observable range	$>$60 km
Antenna gain	38.5 dBi
Antenna diameter	1.2 m
3-dB beamwidth	1.8°
Polarization	Dual linear
Polarimetric isolation	26 dB
Maximum angular velocity	50° s^{-1}
Peak power	100 W
Pulse width	1–69 μs
Chirp bandwidth	4 MHz
Maximum duty cycle	20%
Minimum gate spacing	30 m

solid-state transmitters (one for each polarization), each operating at 100-W peak power. To achieve the desired sensitivity for weather observations despite such low transmit power, pulse compression must be used. A chirp bandwidth of 4 MHz is available for waveform design, and the maximum pulse length is 69 μs, yielding a maximum time-bandwidth product of 276. Additionally, a coupled transmit waveform is recorded on each channel for use as a matched filter. This feature is critical for determining waveform performance and the application of pre-distortion.

The system is transportable and is equipped with mobile Internet for operation from remote locations. An in-house-developed software platform is capable of fully operating the system, ranging from waveform selection to scanning strategy to data management. Raw time series data are available for experimental advanced signal processing; however, moment data are automatically generated and available for simple viewing.

While pulse compression greatly improves the range resolution from such a long pulse (up to 69 μs), we are left with an undesirable side effect: the blind range. While the system is transmitting, no receive data can be collected, meaning that a large unobservable area exists around the radar. For a 69-μs pulse, this equates to a circle of radius 10.35 km, or about one-sixth of the total observable range of the system.

To combat this issue, the method described in Cheong et al. (2013) is used. Instead of utilizing the full 69 μs and 4-MHz bandwidth available for a long pulse, only 67 μs and 2.2 MHz are used for the long pulse. The remaining time and bandwidth are used for an adjacent short pulse, which is called the *fill pulse*. These two subwaveforms are combined into one time–frequency multiplexed transmit waveform that utilizes all available pulse length and bandwidth, and are separated in processing using different matched filters for different ranges. The area within the blind range from the long pulse is filled with data from the short pulse. While sensitivity is low with a short pulse and low transmit power, the distance covered in the blind range is sufficiently small for acceptable sensitivity.

4. Results

While the technique described previously is capable of designing high-performance waveforms in theory, even for low time-bandwidth product radar systems, there are a number of additional steps that must be taken in order to design a waveform for use in a real-world system. The following two sections describe the waveform design from a purely theoretical standpoint, and the changes applied in order to account for transmitter and system distortions. The final section presents actual weather observations using the final waveform design.

a. Theoretical waveform design

In theory, we are able to design a high-performance waveform for the PX-1000 system with two-way SNR loss of 0.05 dB, which represents a large increase over the windowed waveform in Fig. 2a. In a real system, however, there are typically undesired edge characteristics in the transmitter. This means that large spikes of distortion typically occur at the extreme edges of the pulse, which must be mitigated. These major distortions occur most often at both the rising and falling edges of the pulse; however, additional ''droop'' often occurs with long-pulse waveforms in a solid-state transmitter. In the following section, we will explore how to account for most of the distortion in the system; however, these sharp edge distortions are particularly difficult to model. For this reason, the first step to designing a real-world waveform is to start with a slightly windowed pulse in the optimization. We see drastically decreased edge distortion by using a theoretical design window with a 0.24-dB SNR loss. This window is defined as a raised cosine taper on both ends of the waveform with a roll-off factor of 0.1 and assists with the implementation through a real-world solid-state transmitter that contains inherent imperfections.

After the window is applied, the parameters of the system are input into the optimization framework. As noted previously, the specifications of this system include 4-MHz bandwidth and a pulse length of up to 69 μs. However, since the long pulse will be designed separate from the fill pulse, the input into the optimization framework must only utilize the available bandwidth and pulse length for the long pulse. Therefore, the specifications input into the framework for waveform

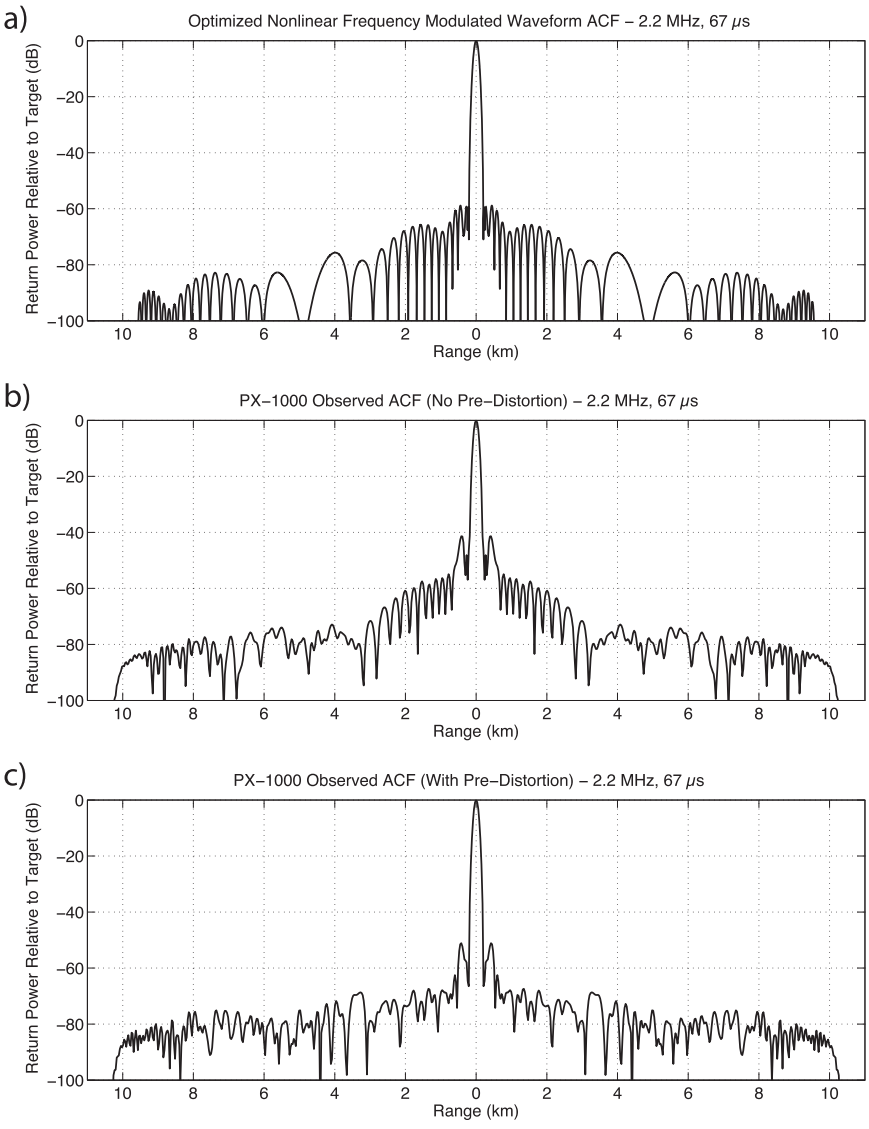

FIG. 5. ACFs of an optimized waveform for the PX-1000 system. (a) The original waveform after optimization with no pre-distortion taken into account. (b) The theoretical waveform from (a) after it has been sent through the transmitter. The ACF is calculated using a coupled pulse. Note the distortion that takes place, causing sidelobes to rise from −59 to −42 dB. (c) The actual coupled ACF of the same optimization technique used for (a) but taking into account pre-distortion. Note that the peak sidelobes have been brought down to −52 dB, a 10-dB increase in performance compared with not accounting for pre-distortion in (b). This is the waveform used in the weather observations in this study.

design in this case are 2.2-MHz bandwidth and 67-μs pulse length, since the fill pulse utilizes the remaining bandwidth and pulse length specifications.

Figure 5a shows the theoretical ACF of the designed waveform for PX-1000 system specifications. A window with two-way SNR loss of 0.24 dB is used, with resulting peak sidelobes of −59 dB, integrated sidelobes of −37 dB, and a 3-dB range resolution of 120 m. While similar sidelobes could be achieved using amplitude modulation with the same system characteristics, the lack of aggressive windowing results in significantly lower power loss as well as high 3-dB range resolution.

A number of other positive aspects of the theoretical waveform are also apparent. In addition to relatively low integrated sidelobes, there is also a consistent drop in sidelobe level. Additionally, as predicted by Bucci and Urkowitz (1993), Griffiths and Vinagre (1994), and George et al. (2008), the Doppler tolerance of the optimized

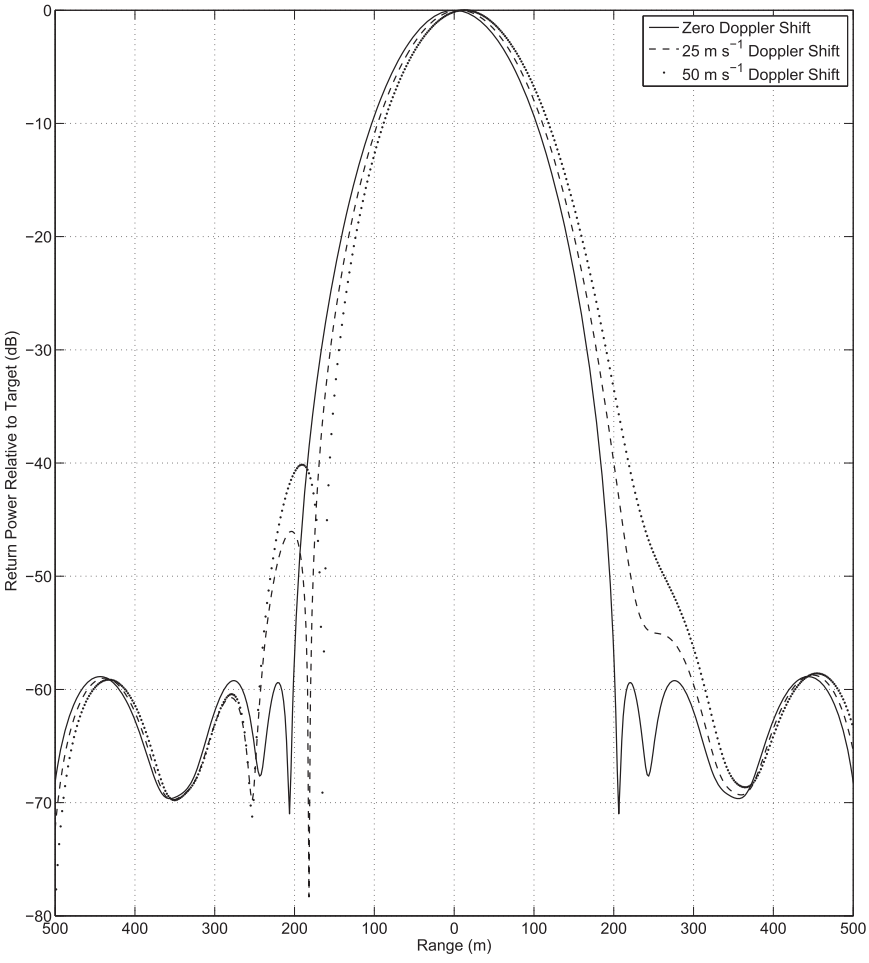

FIG. 6. Cuts through the ambiguity function for a theoretically optimized waveform for the PX-1000 system at 0, 25, and 50 m s^{-1} Doppler shifts. SNR loss of 0.24 dB yielding from a slight raised cosine window is achieved, with peak sidelobes of −59 dB in the ACF at zero Doppler shift. An increase to 50 m s^{-1} in Doppler shift shows little distortion in the main lobe and an increase in peak sidelobe level on one side of the main lobe. Even at higher Doppler shifts, integrated sidelobes outside of the peak sidelobes remain nearly constant, resulting in acceptable overall performance.

waveform is within acceptable range for weather targets. Because of the forcing of a symmetric frequency function, Fig. 6 shows little shift of the main lobe in the ambiguity function with increasing Doppler velocities. The only noticeable side effect at 50 m s^{-1} is an increased peak sidelobe level of −40 dB on one side of the main lobe. For velocities of 25 m s^{-1}, the peak sidelobe level raises to only −46 dB. In each case, the integrated sidelobes remain nearly constant outside of the peak sidelobe, leading to acceptable overall performance in extreme conditions. While this distortion effect at high Doppler velocities may be unavoidable, the symmetric nature of the NLFM design leads to a minimization of error compared with other methods.

In terms of sensitivity gains over more traditional, heavily windowed methods, the optimized waveform displays significant improvements. Figure 7 shows an increase in sensitivity of approximately 2.95 dB throughout the entire range of operation. At 50 km, for example, the optimized NLFM waveform displays a sensitivity under 14 dBZ, while the WLFM lacks observation of echoes below 17 dBZ. This improvement in sensitivity is a critical component in any weather radar platform. It is important to note that there is an abrupt change in sensitivity near the 10-km mark. This is due to the switch between use of the short pulse and the chirped long pulse (the edge of the blind range). For comparison, the popular waveform technique described in De Witte and Griffiths (2004) is also included in Fig. 7. De Witte and Griffiths (2004) utilize a mismatched Nuttall window on receive. The Nuttall window is very aggressive, resulting in an example

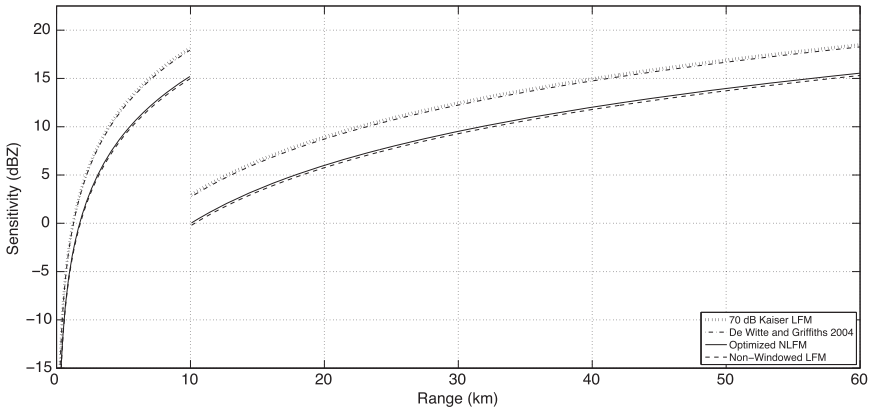

FIG. 7. Theoretical sensitivity comparison between a 70-dB Kaiser window LFM (3.19-dB SNR loss), the Nuttall window mismatched waveform described in De Witte and Griffiths (2004) (2.96-dB SNR loss), the optimized NLFM (0.24-dB SNR loss), and a rectangular LFM (0.00-dB SNR loss). A sensitivity increase of 2.95 dB from the WLFM to the optimized NLFM is observed throughout the operating range. Additionally, the optimized NLFM theoretical sensitivity is within 0.24 dB of the rectangular LFM, indicating very high power efficiency. Note, lower values indicate higher sensitivity, and the first 10 km display poor sensitivity due to the short pulse used to fill the blind range.

where even a mismatched filter can have significant two-way SNR loss (in this case, 2.96 dB compared to 3.19 dB with the 70-dB Kaiser window matched filter LFM). The theoretical sensitivities in Fig. 7 are based on the weather radar equation, from which the reflectivity factor can be calculated (Doviak and Zrnić 1993) using

$$Z = \frac{r^2 l^2 (4\pi)^3 16 \ln 2 P_r \lambda^2}{P_t g^2 c \tau \pi^6 \theta^2 |k_w|^2}, \tag{3}$$

where P_r and P_t are power returned and transmitted, respectively (W); g is the gain of the antenna and l represents losses (both in linear units); c is the speed of light ($\mathrm{m\,s^{-1}}$); τ is the pulse width (s); λ is wavelength and r is range (both in m); θ is one-way half-power beamwidth; and k_w is the dielectric factor of water at a specified temperature (unitless). Power returned P_r is set to the noise floor of the system, which is equal to -110 dBm for the PX-1000 platform.

The use of a matched filter in waveform design and processing is supported by various results in the literature, most notably those contained within Harris (1978). While a mismatched filter, even with lower bandwidth on receive, would result in a lower noise floor because of lower receiver noise, Harris (1978) shows that no combination of mismatched filters can overcome the gain in sensitivity afforded by using a matched filter. While there is no separate theorem for distributed targets such as hydrometeors, system tests with PX-1000 consistently verify this concept for weather observations.

A comparison of the theoretical power spectrum between a nonwindowed LFM, a 70-dB Kaiser window LFM, and the optimized NLFM, all at 2.2-MHz chirp bandwidth, is provided in Fig. 8a. The optimized waveform, due to the lack of windowing, shows a slight broadening in the spectrum above -75 dB compared with the heavily windowed waveform. However, the optimized waveform utilizes distinctively less spectrum than the nonwindowed LFM. It is important to note that from a spectral efficiency standpoint, the use of a slightly windowed, matched filter waveform will use less overall spectrum than a rectangular transmit waveform with a mismatched receive filter, despite the slight increase in receiver noise.

Finally, in order to measure the effects of system phase noise on the waveform technique, a series of 100 phase noise simulations was averaged in Fig. 8b. As expected, an increase in phase noise results in decreased sidelobe performance, similar to the results shown in Bharadwaj and Chandrasekar (2012). The estimated phase noise of PX-1000, however, is only approximately 0.05°, which is marked with a black diamond in Fig. 8b. At this level, the effect on waveform PSL performance is approximately 0.08 dB. Systems with higher phase noise will show decreased performance as phase noise increases.

b. *Accounting for transmitter distortion*

As previously mentioned, transmitter nonlinearities lead to a transmitted waveform that can be significantly altered from what was intended. Despite the ideal theoretical waveform ACF shown in Fig. 5a, the waveform actually transmitted by PX-1000 results in a much less desirable ACF. Figure 5b shows the ACF of the coupled pulse after sending the theoretically optimal waveform

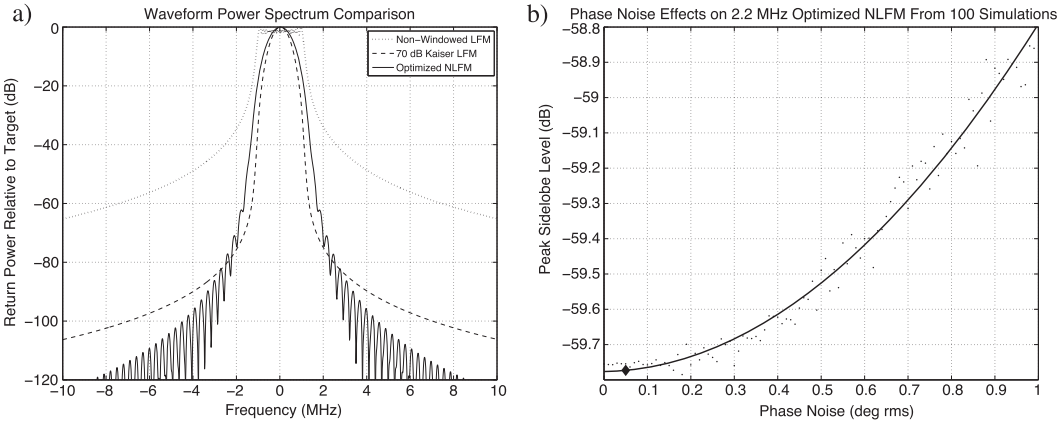

FIG. 8. (a) Theoretical power spectrum comparison between a rectangular LFM, 70-dB Kaiser window LFM, and the optimized NLFM, all with 2.2-MHz chirp bandwidth. The lack of windowing on the optimized waveform causes a slight broadening compared with the windowed LFM in the spectrum above −75 dB. (b) Estimate of phase noise effects on the 2.2-MHz theoretical optimized NLFM waveform using an average of 100 simulations. A polynomial fit is shown via the thick black line, and the system phase noise estimate of PX-1000 is marked with a black diamond (0.05°), showing no appreciable detrimental effects in the system.

through the transmitter. It is clear that performance has been significantly degraded, and the transmitter results in a number of undesirable effects.

The measured coupled pulse results in peak sidelobes of −42 dB, integrated sidelobes of −26 dB, and a 3-dB range resolution of 120 m. The SNR loss has remained at 0.24 dB. With a theoretical waveform design resulting in peak and integrated sidelobes of −59 and −37 dB, respectively, the distortion clearly leads to a waveform that is not acceptable. To mitigate this issue, we use a measured inverse transfer function in the optimization process (see Fig. 4b) to account for transmitter distortion.

While a slightly more aggressive window with a 0.24-dB SNR loss, rather than a window with 0.05-dB SNR loss, was used in the original theoretical design in order to lessen noisy transmitter edge effects, there are still distortions that occur away from the edges of the pulse. Both frequency and amplitude distortions occur, with the most obvious being a droop in amplitude from the beginning of the pulse to the end. This is where the inverse transfer function accounts for distortions accurately and can drastically improve performance.

Figure 5c shows the coupled result of the waveform ACF from Fig. 5a optimized with pre-distortion taken into account. Figure 5c can be directly compared to Fig. 5b to see improvement of the waveform ACF through the actual system. Actual observed sidelobes have decreased from −42 to −52 dB, and integrated sidelobes have decreased from −26 to −33 dB. Two-way SNR loss has remained at 0.24 dB, and the 3-dB range resolution has remained at 120 m. The results are significantly closer to theoretical values.

c. *Weather observations*

To demonstrate the increase in sensitivity afforded over heavy windowing methods, this section presents three weather cases observed with the PX-1000 system. The key aspects to waveform performance, specifically for weather observations, are low sidelobes and high sensitivity. In each of the following cases, the optimized NLFM waveform (OFM) excels in both of these aspects. It is capable of retaining nearly identical sensitivity to a nonwindowed LFM waveform, but displays no increase in sidelobes over a more traditional WLFM. It is important to note that this is a small collection of cases that demonstrates high waveform performance for the OFM. Marked increases in performance of all weather radar moments and parameters have been observed in all cases that have been collected. In each case, a comparison of results using a nonwindowed LFM waveform, a WLFM waveform, and the OFM waveform, all with PX-1000, are compared with the estimated moments from a WSR-88D radar (either KTLX or KOUN). In case 3, specifically, the OFM also includes the use of a time–frequency multiplexed fill pulse in order to assist in blind range mitigation (Cheong et al. 2013).

1) CASE 1: HIGH REFLECTIVITY GRADIENT, 9 JULY 2012

Figure 9 shows the classic situation where a strong reflectivity gradient can result in undesirable sidelobes. In Fig. 9a, an LFM waveform used to observe a strong reflectivity gradient results in an extension of reflectivity toward the radar. Not only does this distort weather targets but the artificial observations are not corrected in

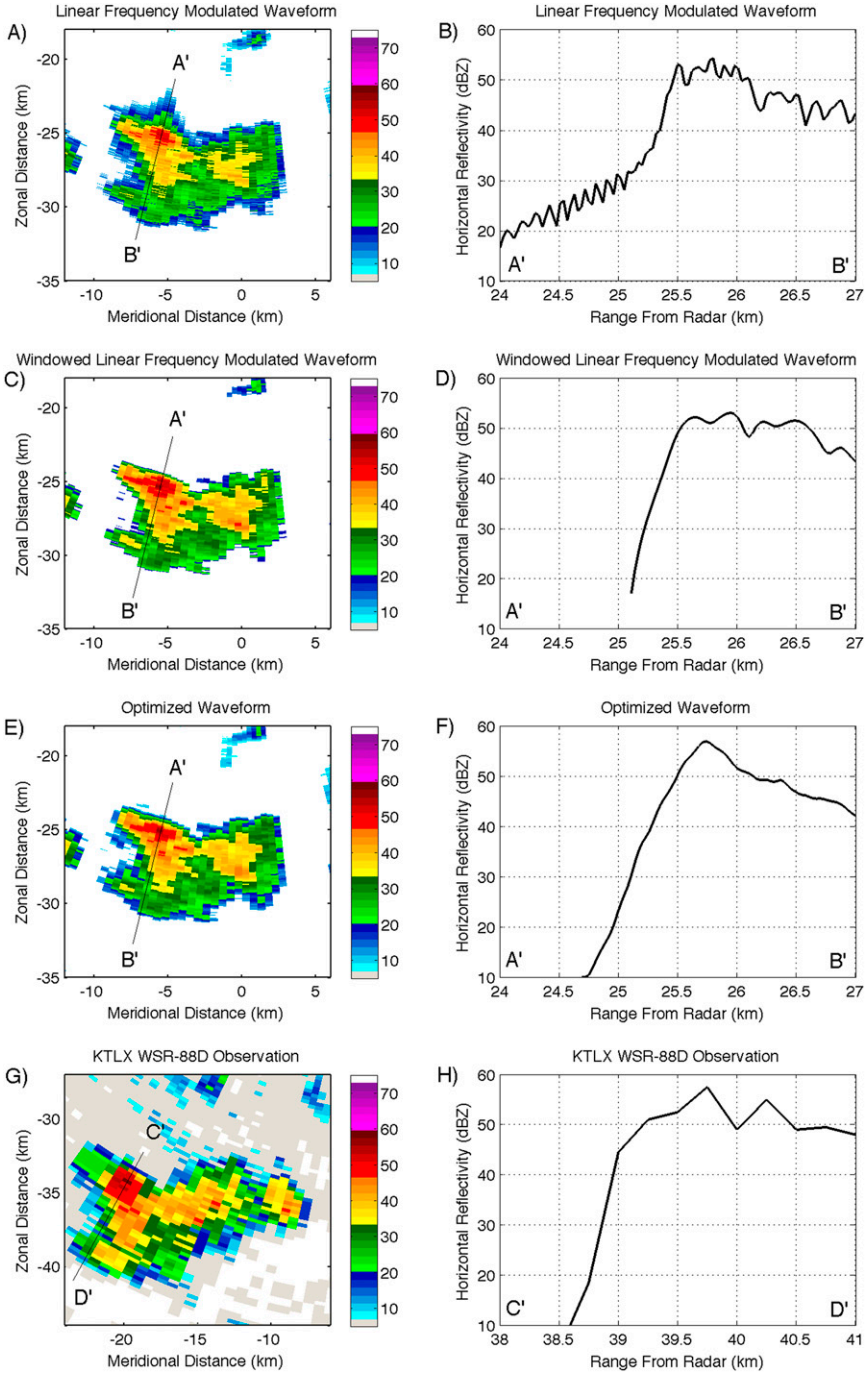

FIG. 9. Observations of a convective cell with a strong reflectivity gradient at 2028 UTC 9 Jul 2012 in Norman, OK. The left-hand side shows reflectivity (dBZ) in a planned position indicator, and the right-hand side shows a normalized range profile view through the highest area of reflectivity, simulating an ACF and displaying sidelobes. In the range profile plots, the left-hand side indicates range toward the northeast (toward the radar), while the right-hand side indicates range toward the southwest (away from the radar). Collected with (a),(b) an LFM waveform, (c),(d) a WLFM waveform, (e),(f) an OFM waveform, and (g),(h) the KTLX WSR-88D at approximately the same time.

derived estimates. This means that rainfall estimations, for example, will appear to be anomalously high. In addition to strong reflectivity gradients, sidelobes can contaminate clustered areas of cells without knowledge of the user. Figure 9b shows a range profile through the area of highest reflectivity, with the left side pointing toward the radar (to the northeast) and the right side pointing away from the radar (to the southwest). It is clear that high sidelobes are contaminating the signal to the northeast of the convection.

Figure 9c shows the same convective cell approximately 20 s later. This time, a WLFM waveform is used with an SNR loss of 3.19 dB. While the artificial extension in reflectivity has disappeared because of high sidelobe suppression, a marked loss in sensitivity is observed. All reflectivity values below approximately 15 dBZ are no longer observable because of the heavy windowing of the pulse. This alters the shape and structure of the cell as it appears to the user, with similar negative impacts to those discussed in the LFM case. Figure 9d shows that the high sidelobes from Fig. 9b have been suppressed, but at a cost to sensitivity. It is also important to note that data with an SNR value below 0 are filtered out in all of the PX-1000 cases shown.

Figure 9e shows the same cell approximately 20 s after Fig. 9c. In this case, an OFM waveform is used, with an SNR loss of 0.24 dB. In addition to no artificial sidelobes, the sensitivities observed in the LFM case have been recovered, showing accurate edge characteristics for the cell in question. Figure 9f shows that no sidelobes are being observed and that we have an accurate representation of the cell.

Finally, Figs. 9g,h show the same cell observed at approximately the same time by the KTLX WSR-88D at Twin Lakes, Oklahoma. Despite being farther away from the storm, the same sharp reflectivity gradient is seen. Because of the significantly higher power of the WSR-88D system, additional sensitivity below 10 dBZ is possible; however, these values are below those capable of being observed with a 100-W radar platform, even with a long pulse. Regardless of the power difference, the same general storm shape is seen, including the important outer edges seen in both the LFM and OFM waveforms. Table 2 shows the relevant system characteristics of the WSR-88D platform for comparison to PX-1000.

TABLE 2. System characteristics of WSR-88D (Doviak et al. 2000).

Transmitter type	Klystron
Operating frequency	2700–3000 MHz
Sensitivity	−27 dBZ at 50 km
Observable range	460 km
Antenna gain	45.5 dBi
Antenna diameter	8.5 m
3-dB beamwidth	0.95°
Polarization	Simultaneous linear
Polarimetric isolation	37 dB
Maximum angular velocity	36° s^{-1}
Peak power	750 kW
Pulse width	1.57–4.5 μs
System bandwidth	0.63 MHz
Maximum duty cycle	2%
Minimum gate spacing	250 m (in super-resolution mode)

2) CASE 2: STRATIFORM/CONVECTIVE MULTICELLS, 18 AUGUST 2012

Waveform performance is critical in all weather situations, not just those that involve sharp reflectivity gradients. Additionally, it is important to be sure that other estimates aside from reflectivity agree with those from nearby radar systems. Figure 10 shows a case involving a mix of stratiform and convective precipitation in the summer months. Despite the lack of clear reflectivity gradients, changes in sensitivity are apparent between different waveforms. The WLFM waveform lacks returns below 20 dBZ at ranges greater than 25 km from the radar. This is most apparent at the hole in reflectivity in the northwest quadrant of each plot. While the LFM waveform partially fills the hole with lower reflectivity values, these values are not observable with the WLFM. Because no return is detected, which is indicated by low SNR and low spectrum width in the low reflectivity area, other estimations are not possible. Velocity estimates and dual-polarimetric estimates are nonexistent in this area using the WLFM compared with the LFM.

The third column in Fig. 10 shows the same scenario approximately 20 s later but observed with the OFM waveform. The gap in low signal return has been mostly filled in around the edges with lower reflectivity values. The fourth column was collected at approximately the same time by the collocated KOUN WSR-88D. It is clear that the area of low SNR should be mostly filled with low reflectivity values, indicating that the OFM has retained significantly more sensitivity than the WLFM. While sidelobes are not readily apparent in the LFM waveform in this case, it should be noted that without sharp reflectivity gradients it is difficult to observe sidelobes. Simply via the nature of nonwindowed LFM waveforms, however, it must be assumed that the data are corrupted by sidelobes.

In this case, where sufficient sensitivity is retained, velocity estimates are not affected by the use of pulse compression waveforms. As long as a signal is received, velocity estimates appear to closely match the WSR-88D estimates. To cover 60 km, the PX-1000 configuration had an aliasing velocity of 15 m s^{-1}. To correct for this, a simple dealiasing algorithm was applied to the PX-1000 observations.

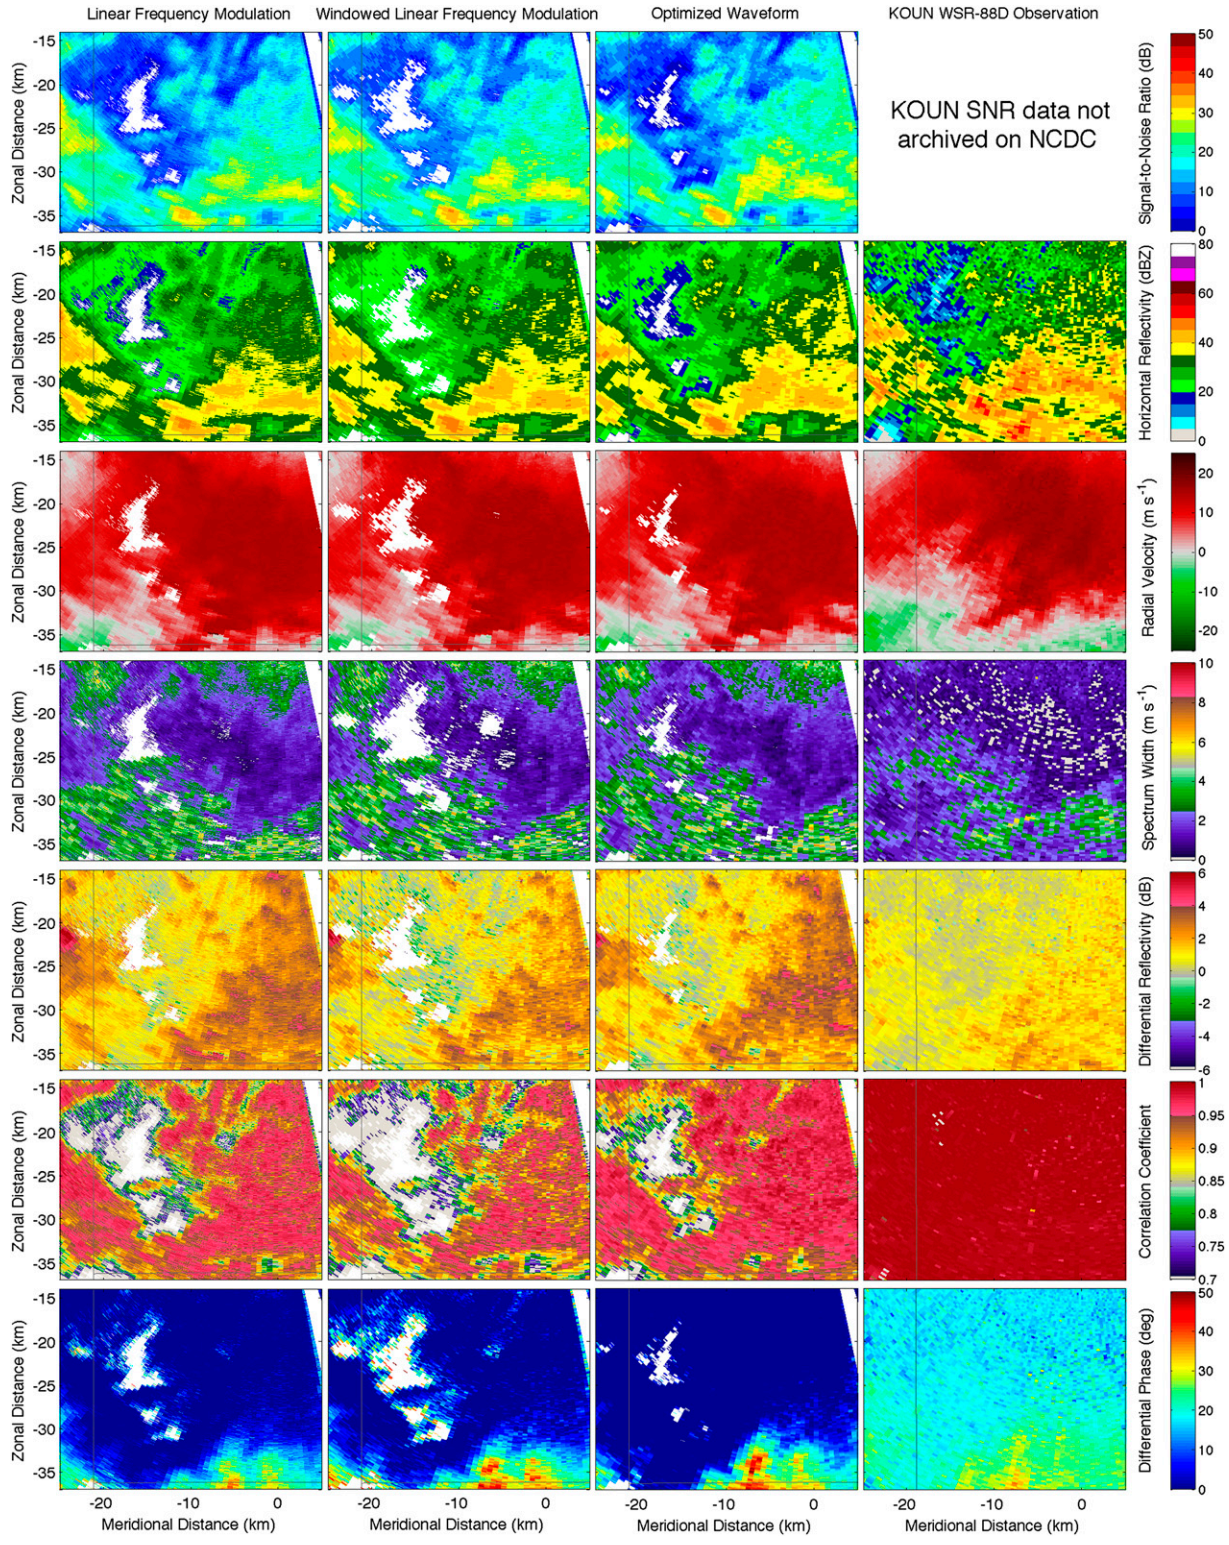

FIG. 10. Observations of a mixed stratiform/convective multicellualar cluster at 1914 UTC 18 Aug 2012 in Norman. The columns represent collection with (from left to right) an LFM waveform, a WLFM waveform, an OFM waveform, and the collocated KOUN WSR-88D, at approximately the same time. The rows show (from top to bottom) SNR (dB), horizontal reflectivity (dBZ), radial velocity ($\mathrm{m\,s^{-1}}$), spectrum width ($\mathrm{m\,s^{-1}}$), differential reflectivity (dB), correlation coefficient, and differential phase (°).

The bottom three rows in Fig. 10 show dual-polarization estimates, which are an important aspect of any weather radar platform. It is important for pulse compression waveforms to deliver similar polarization performance as their short-pulse counterparts. A key aspect to dual-polarization moment estimation with pulse compression is the retainment of sensitivity. As sensitivity falls off, error in dual-polarization moment estimations rises considerably (Lei et al. 2012). This is apparent primarily in large areas of nonnatural correlation coefficient values observed with the WLFM waveform in areas with low SNR. Because of the loss in sensitivity from windowing, relatively few areas in the precipitation display expected correlation coefficient values higher than 0.95. While the OFM certainly does not match the KOUN observations, a significant recovery is evident over the WLFM waveform. Low-powered radars, even with high-power efficiency pulse compression waveforms, are an example where the multilag method described by Lei et al. (2012) could be used for more accurate correlation coefficient values. Such an approach would only be strengthened by using an OFM waveform compared to a WLFM waveform.

Similar performance can be seen in differential reflectivity. It is important to note that differences in wavelength can lead to different polarimetric estimates (Ryzhkov and Zrnić 2005; Ryzhkov 2007; Kumjian and Ryzhkov 2008). In this case, we see similar spatial patterns in all radar moment estimates, and the differences we would expect between S band and X band are apparent. The only estimates that do not necessarily match the KOUN observations are the reflectivity estimates. Because of the use of a low-power X-band transmitter, even a long pulse cannot avoid attenuation in convective precipitation. As the transmitted signal travels farther through precipitation, signal loss becomes increasingly apparent. As long as some signal is returned, however, attenuation correction methods can be applied, given that differential phase estimates are accurate (Bringi et al. 1990; Gorgucci and Chandrasekar 2005; Park et al. 2005; Snyder et al. 2010). The bottom row in Fig. 10 shows reasonable differential phase values at X band, leading to attenuation-corrected values that come close to KOUN estimates.

3) Case 3: Convective line segments, 26 August 2012

One of the primary concerns for long-pulse radar systems is the blind range. While multiple studies have attempted to mitigate the blind range (most notably Bharadwaj and Chandrasekar 2012; Cheong et al. 2013), the method described in Cheong et al. (2013) is being tested on PX-1000. Figure 11 presents a situation with two convective line segments: one to the west of the radar and one directly over and to the east of the radar. It is clear that in this situation, seen in full detail in the fourth column as observed by the collocated KOUN WSR-88D radar, a solution must be developed in order to see both convective lines.

While Fig. 11 shows the same method for collecting pulse compression data with the LFM and WLFM as before, the third column shows an example of a time–frequency multiplexed OFM waveform. After a long pulse of 67 μs is transmitted, a short pulse of 2 μs is transmitted at the end of the waveform. Areas outside of the blind range are match filtered with the long pulse for maximum sensitivity, while areas within the blind range are match filtered with the short pulse in order to fill in the area that was invisible to the radar while transmitting the long pulse. Using this method combined with an OFM waveform, the third column in Fig. 11 shows estimates both outside and inside the blind range. While not fully resolved because of the inherent lack of sensitivity with a short pulse and 100-W peak transmit power within the blind range, the leading convective line that passes over the radar is observed. As with case 2, an application of attenuation correction can be useful in situations where reflectivity bias correction is needed.

In the other moment estimates, similar results to those noted in case 2 are observed. The OFM waveform shows a considerable increase in sensitivity over the WLFM, as is evident by increased SNR, especially to the east of the radar. This is a particularly critical area for power efficiency and radar sensitivity in this example, since an area of heavy convection is directly over and to the east of the radar. We see a significant increase in sensitivity in the OFM waveform to the east, as returns between 10 and 15 dBZ are observable. As with case 2, the increased sensitivity afforded by the OFM waveform allows for more accurate estimates of all moments, and estimates of moments that were not observable with traditional pulse compression methods.

5. Conclusions and future work

This paper presents a new approach to designing weather radar waveforms for pulse compression. As opposed to previous methods, which often relied upon aggressive amplitude modulation with low power efficiencies, our approach yields promising results with theoretical SNR losses as low as 0.05 dB, and actual SNR losses in practice as low as 0.24 dB. By using a flexible Bézier curve method for the frequency function and an appropriately weighted cost function within a genetic optimization framework, low peak and integrated sidelobes, as well as high range resolution, are possible. The results in this study were based on waveform designs for a 2.2-MHz bandwidth, 67-μs pulse length weather radar system with only 100 W of peak transmit power on each channel.

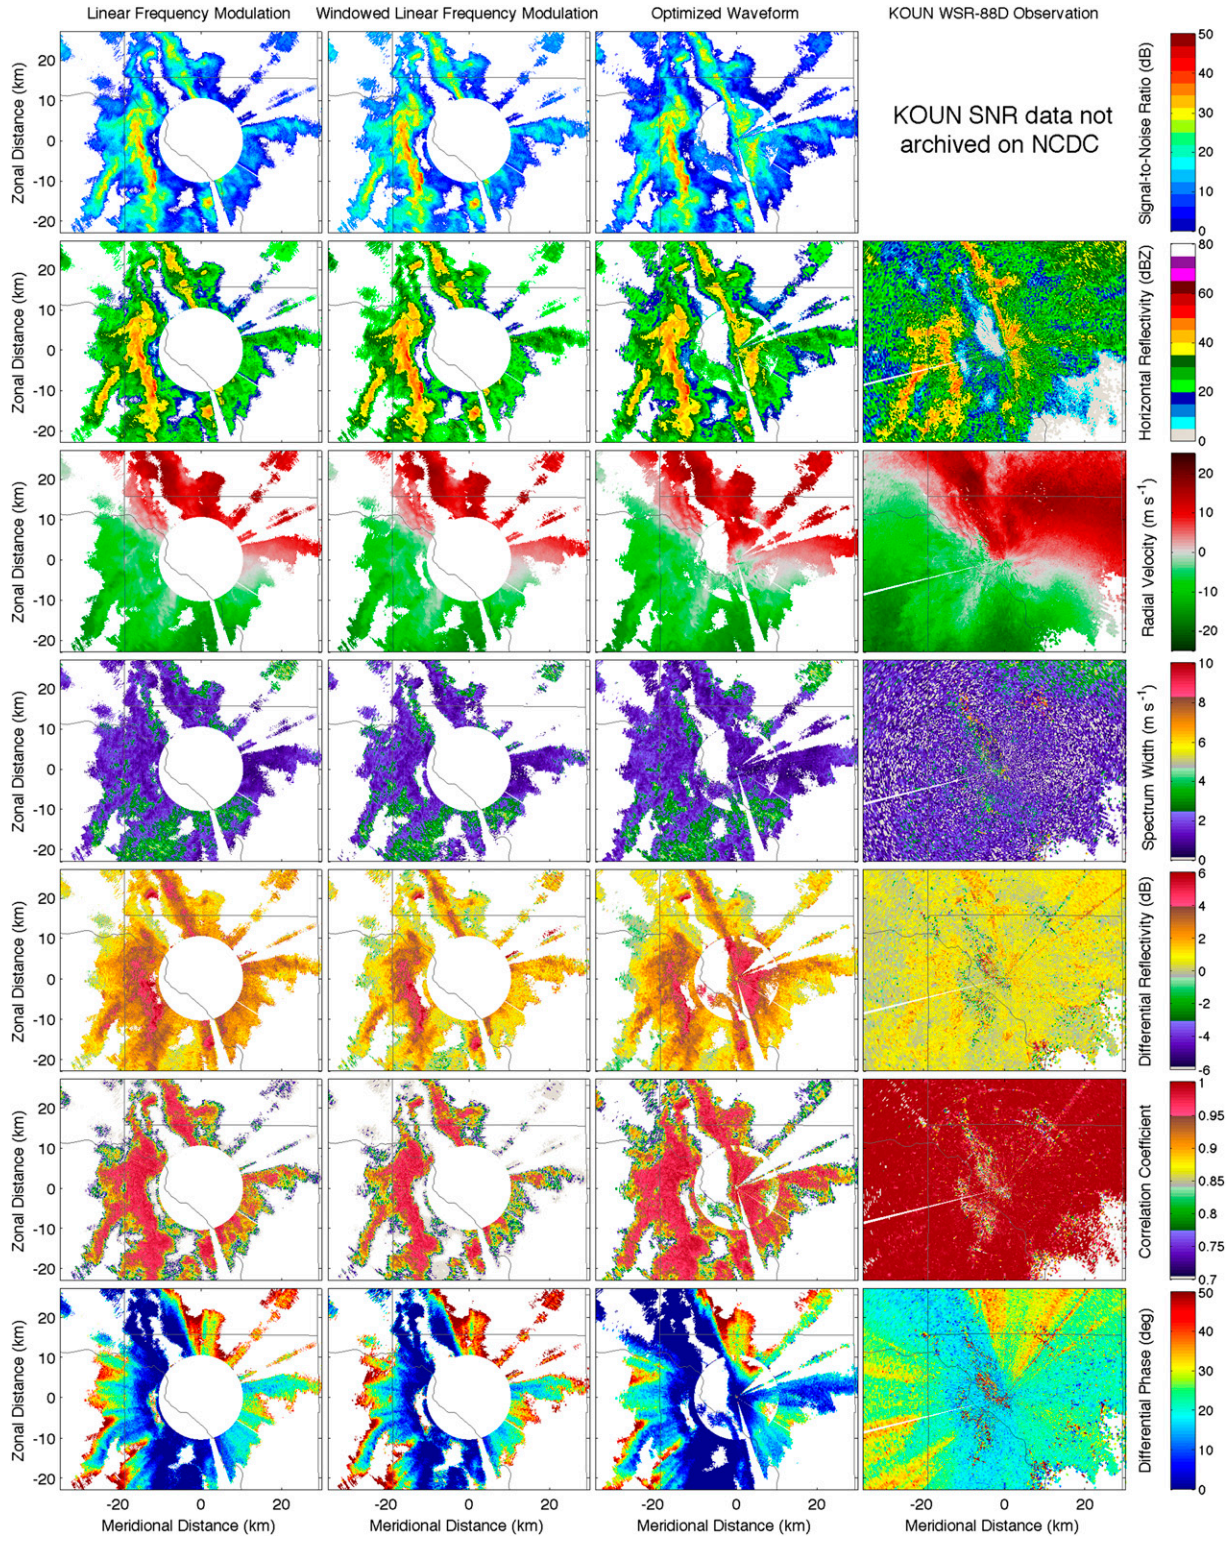

FIG. 11. Observations of two convective line segments at 0453 UTC 26 Aug 2012 in Norman. The columns represent collection with (from left to right) an LFM waveform, a WLFM waveform, an OFM waveform, and the collocated KOUN WSR-88D, at approximately the same time. The rows show (from top to bottom) SNR (dB), horizontal reflectivity (dBZ), radial velocity ($m\,s^{-1}$), spectrum width ($m\,s^{-1}$), differential reflectivity (dB), correlation coefficient, and differential phase (°).

In addition to designing theoretical waveforms that are capable of high-performance pulse compression results, this study also utilizes a method to account for transmitter distortions when implementing designed waveforms in real-world systems. By measuring a transfer function of the transmit chain and applying the inverse transfer function within the optimization framework, a significant improvement in performance is observed in implementation. On the PX-1000 system, an improvement from −42 to −52 dB was observed for peak sidelobes, and an improvement from −26 to −32 dB was observed for integrated sidelobes.

To efficiently utilize such a low-power transmitter in a weather radar platform, pulse compression is a necessity; 100 W is not sufficient to observe weak weather targets with a short pulse, and with traditional pulse compression methods, a low power efficiency is also not ideal. Gain in sensitivity as high as 2.95 dB is observed over traditional, aggressively windowed methods. This is significant in the realm of weather radar because of the nature of observations that are desired. A higher dynamic range of weather targets are now observable with minimized range sidelobes.

In addition to application in numerous traditional, inexpensive, low-power dish radar systems, pulse compression methods such as this for weather radar have significant potential for future phased array systems. Because of the desire for features such as high temporal resolution, the lack of moving parts, lower maintenance costs, and graceful degradation, low-power transmitters will be important in future phased array radar designs, both for weather and military uses. To keep costs reasonable, low-power transmitters and effective pulse compression waveforms will be critical. The methods presented in this study provide a viable option for simple pulse compression implementation in solid-state transmitters, resulting in accurate estimations and few side effects.

Numerous topics exist for future work in this area. Theoretical waveforms have been developed for much higher time-bandwidth systems, but have not been tested with predistortion on actual systems. Additionally, some work has been completed implementing this technique on TWT transmitters. TWT transmitters tend to have more difficult edge effects to work with and therefore lead to higher distortion rates. The inverse transfer function method does not work as well in these situations and may need refinement. Future work with waveform development for the University of Oklahoma's cylindrical polarimetric phased array radar (CPPAR) is anticipated in the near future (Zhang et al. 2011) along with additional optimization applications, taking into account spectrum usage and Doppler tolerance.

Acknowledgments. This work was supported by the National Severe Storms Laboratory (NOAA/NSSL) under the Cooperative Agreement NA08OAR4320904. The authors thank Redmond Kelley for the numerous discussions regarding waveform design and system implementation. We also thank the Toshiba Corporation for the useful discussions regarding waveform techniques, as well as the opportunity to further test our method on their radar systems. Finally, the authors are grateful for the comments provided by the three anonymous reviewers, which significantly increased the technical reach of the paper.

REFERENCES

Ackroyd, M. H., and F. Ghani, 1973: Optimum mismatched filters for sidelobe suppression. *IEEE Trans. Aerosp. Electron. Syst.,* **AES-9,** 214–218, doi:10.1109/TAES.1973.309769.

Austin, G. L., 1974: Pulse compression systems for use with meteorological radars. *Radio Sci.,* **9,** 29–33, doi:10.1029/RS009i001p00029.

Baden, J. M., and M. N. Cohen, 1990: Optimal peak sidelobe filters for biphase pulse compression. *Record of the IEEE 1990 International Radar Conference,* IEEE, 249–252.

Bharadwaj, N., and V. Chandrasekar, 2012: Wideband waveform design principles for solid-state weather radars. *J. Atmos. Oceanic Technol.,* **29,** 14–31, doi:10.1175/JTECH-D-11-00030.1.

Bluestein, H. B., M. M. French, I. PopStefanija, R. T. Bluth, and J. B. Knorr, 2010: A mobile, phased-array Doppler radar for the study of severe convective storms: The MWR-05XP. *Bull. Amer. Meteor. Soc.,* **91,** 579–600, doi:10.1175/2009BAMS2914.1.

Bringi, V. N., V. Chandrasekar, N. Balakrishnan, and D. S. Zrnić, 1990: An examination of propagation effects in rainfall on radar measurements at microwave frequencies. *J. Atmos. Oceanic Technol.,* **7,** 829–840, doi:10.1175/1520-0426(1990)007<0829:AEOPEI>2.0.CO;2.

Brown, R. A., and V. T. Wood, 2012: Simulated vortex detection using a four-face phased-array Doppler radar. *Wea. Forecasting,* **27,** 1598–1603, doi:10.1175/WAF-D-12-00059.1.

Bucci, N. J., and H. Urkowitz, 1993: Testing of Doppler tolerant range sidelobe suppression in pulse compression meteorological radar. *Record of the 1993 IEEE National Radar Conference,* IEEE, 206–211.

——, H. S. Owen, K. A. Woodward, and C. M. Hawes, 1997: Validation of pulse compression techniques for meteorological functions. *IEEE Trans. Geosci. Remote Sens. Lett.,* **35,** 507–523, doi:10.1109/36.581958.

Carswell, J., S. Bidwell, and R. Meneghini, 2008: A novel solid-state, dual-polarized, dual wavelength precipitation Doppler radar/radiometer. *2008 IEEE International Geoscience and Remote Sensing Symposium: Proceedings,* Vol. 5, IEEE, IV-1014–IV-1017.

Cheong, B. L., R. Kelley, R. D. Palmer, Y. Zhang, M. Yeary, and T.-Y. Yu, 2013: PX-1000: A solid-state polarimetric X-band radar and time-frequency multiplexed waveform for blind range mitigation. *IEEE Trans. Instrum. Meas.,* **62,** 3064–3072, doi:10.1109/TIM.2013.2270046.

Cilliers, J. E., and J. C. Smit, 2007: Pulse compression sidelobe reduction by minimization of L_p-norms. *IEEE Trans. Aerosp. Electron. Syst.,* **43,** 1238–1247, doi:10.1109/TAES.2007.4383616.

De Witte, E., and H. D. Griffiths, 2004: Improved ultra-low range sidelobe pulse compression waveform design. *IEEE Electron. Lett.,* **40,** 1448–1450, doi:10.1049/el:20046548.

Doviak, R., and D. Zrnić, 1993: *Doppler Radar and Weather Observations.* 2nd ed. Academic Press, 562 pp.

Doviak, R. J., V. Bringi, A. Ryzhkov, A. Zahrai, and D. Zrnić, 2000: Considerations for polarimetric upgrades to operational WSR-88D radars. *J. Atmos. Oceanic Technol.,* **17,** 257–278, doi:10.1175/1520-0426(2000)017<0257:CFPUTO>2.0.CO;2.

Duh, F. B., C. F. Juang, and C. T. Lin, 2004: A neural fuzzy network approach to radar pulse compression. *IEEE Geosci. Remote Sens. Lett.,* **1,** 15–20, doi:10.1109/LGRS.2003.822310.

Eiben, A. E., and J. E. Smith, 2007: *Introduction to Evolutionary Computing.* 2nd ed. Natural Computing Series, Springer, 300 pp.

Farin, G. E., 1997: *Curves and Surfaces for CAGD: A Practical Guide.* 4th ed. Computer Science and Scientific Computing, Academic Press, 429 pp.

Ge, Z., P. Huang, and W. Lu, 2008: Matched NLFM pulse compression method with ultra-low sidelobes. *EuRAD 2008: European Radar Conference,* IEEE, 92–95.

George, J., N. Bharadwaj, and V. Chandrasekar, 2008: Considerations in pulse compression design for weather radars. *2008 IEEE International Geoscience and Remote Sensing Symposium: Proceedings,* Vol. 5, IEEE, V-109–V-112.

Gorgucci, E., and V. Chandrasekar, 2005: Evaluation of attenuation correction methodology for dual-polarization radars: Application to X-band systems. *J. Atmos. Oceanic Technol.,* **22,** 1195–1206, doi:10.1175/JTECH1763.1.

Griffiths, H. D., and L. Vinagre, 1994: Design of low-sidelobe pulse compression waveforms. *IEEE Electron. Lett.,* **30,** 1004–1005, doi:10.1049/el:19940644.

Harris, F. J., 1978: On the use of windows for harmonic analysis with the discrete Fourier transform. *Proc. IEEE,* **66,** 51–83, doi:10.1109/PROC.1978.10837.

Heinselman, P. L., D. L. Priegnitz, K. L. Manross, T. M. Smith, and R. W. Adams, 2008: Rapid sampling of severe storms by the National Weather Radar Testbed Phased Array Radar. *Wea. Forecasting,* **23,** 808–824, doi:10.1175/2008WAF2007071.1.

Hirth, B. D., J. L. Schroeder, S. W. Gunter, and J. G. Guynes, 2012: Measuring a utility-scale turbine wake using the TTUKa mobile research radars. *J. Atmos. Oceanic Technol.,* **29,** 765–771, doi:10.1175/JTECH-D-12-00039.1.

Holland, J., 1975: *Adaption in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence.* University of Michigan Press, 183 pp.

Keeler, R. J., and C. A. Hwang, 1995: Pulse compression for weather radar. *Record of the IEEE 1995 International Radar Conference,* IEEE, 529–535.

Kumjian, M. R., and A. V. Ryzhkov, 2008: Polarimetric signatures in supercell thunderstorms. *J. Appl. Meteor. Climatol.,* **47,** 1940–1961, doi:10.1175/2007JAMC1874.1.

Lei, L., G. Zhang, R. J. Doviak, R. Palmer, B. L. Cheong, M. Xue, Q. Cao, and Y. Li, 2012: Multilag correlation estimators for polarimetric radar measurements in the presence of noise. *J. Atmos. Oceanic Technol.,* **29,** 772–795, doi:10.1175/JTECH-D-11-00010.1.

Levanon, N., and E. Mozeson, 2004: *Radar Signals.* Wiley, 432 pp.

McLaughlin, D., and Coauthors, 2009: Short-wavelength technology and the potential for distributed networks of small radar systems. *Bull. Amer. Meteor. Soc.,* **90,** 1797–1817, doi:10.1175/2009BAMS2507.1.

Mudukutore, A. S., V. Chandrasekar, and R. J. Keeler, 1998: Pulse compression for weather radars. *IEEE Trans. Geosci. Remote Sens.,* **36,** 125–142, doi:10.1109/36.655323.

O'Hora, F., and J. Bech, 2007: Improving weather radar observations using pulse-compression techniques. *Meteor. Appl.,* **14,** 389–401, doi:10.1002/met.38.

Park, S. G., V. N. Bringi, V. Chandrasekar, M. Maki, and K. Iwanami, 2005: Correction of radar reflectivity and differential reflectivity for rain attenuation at X band. Part I: Theoretical and empirical basis. *J. Atmos. Oceanic Technol.,* **22,** 1621–1632, doi:10.1175/JTECH1803.1.

Ryzhkov, A. V., 2007: The impact of beam broadening on the quality of radar polarimetric data. *J. Atmos. Oceanic Technol.,* **24,** 729–744, doi:10.1175/JTECH2003.1.

——, and D. S. Zrnić, 2005: Radar polarimetry at S, C, and X bands: Comparative analysis and operational implications. *32nd Conf. Radar Meteorology,* Albuquerque, NM, Amer. Meteor. Soc., 9R.3. [Available online at https://ams.confex.com/ams/32Rad11Meso/techprogram/paper_95684.htm.]

Snyder, J. C., H. B. Bluestein, G. Zhang, and S. J. Frasier, 2010: Attenuation correction and hydrometeor classification of high-resolution, X-band, dual-polarized mobile radar measurements in severe convective storms. *J. Atmos. Oceanic Technol.,* **27,** 1979–2001, doi:10.1175/2010JTECHA1356.1.

Wang, P., H. Meng, and W. Xiqin, 2008: Suppressing autocorrelation sidelobes of LFM pulse trains with genetic algorithm. *Tsinghua Sci. Technol.,* **13,** 800–806, doi:10.1016/S1007-0214(08)72203-X.

Weadon, M., P. Heinselman, D. Forsyth, J. Kimpel, W. E. Benner, and G. S. Torok, 2009: Multifunction phased array radar. *Bull. Amer. Meteor. Soc.,* **90,** 385–389, doi:10.1175/2008BAMS2666.1.

Weber, M. E., J. Y. N. Cho, J. S. Herd, J. M. Flavin, W. E. Benner, and G. S. Torok, 2007: The next-generation multimission U.S. surveillance radar network. *Bull. Amer. Meteor. Soc.,* **88,** 1739–1751, doi:10.1175/BAMS-88-11-1739.

Yichun, P., P. Shirui, Y. Kefeng, and D. Wenfeng, 2005: Optimization design of NLFM signal and its pulse compression simulation. *2005 IEEE International Radar Conference,* IEEE, 383–386.

Yussouf, N., and D. J. Stensrud, 2008: Impact of high temporal frequency radar data assimilation on storm-scale NWP model simulations. *24th Conf. on Severe Local Storms,* Savannah, GA, Amer. Meteor. Soc., 9B.1. [Available online at https://ams.confex.com/ams/24SLS/techprogram/paper_141555.htm.]

Zhang, G., R. J. Doviak, D. S. Zrnić, R. Palmer, L. Lei, and Y. Al-Rashid, 2011: Polarimetric phased-array radar for weather measurement: A planar or cylindrical configuration? *J. Atmos. Oceanic Technol.,* **28,** 63–73, doi:10.1175/2010JTECHA1470.1.

Zrnić, D. S., and Coauthors, 2007: Agile-beam phased array radar for weather observations. *Bull. Amer. Meteor. Soc.,* **88,** 1753–1766, doi:10.1175/BAMS-88-11-1753.