P4R.8 ENHANCED RADAR DATA ACQUISITION SYSTEM AND SIGNAL PROCESSING ALGORITHMS FOR THE TERMINAL DOPPLER WEATHER RADAR

John Y. N. Cho*, Gabriel R. Elkin, and Nathan G. Parker
MIT Lincoln Laboratory, Lexington, Massachusetts

1. INTRODUCTION

As part of a broader FAA program to improve supportability, the Terminal Doppler Weather Radar (TDWR) radar data acquisition (RDA) subsystem is being replaced. For this purpose we developed an engineering prototype RDA with a scalable, open-systems hardware platform. This paper describes the design and characteristics of this new system. The dramatically increased computing power and more flexible transmitter control also enables modern signal processing algorithms to be implemented to improve the quality of the base data. Results highlighting the improved range-overlay protection provided by the new algorithms are presented.

2. RDA PROTOTYPE

Figure 1 shows a simplified block diagram of the TDWR as currently installed at 45 U.S. sites. The digital signal processor (DSP) performs clutter filtering and generates moment data, as well as functioning as a conduit for system control between the Remote Monitoring System (RMS) and the antenna, transmitter, and receiver/exciter (REX) subsystems. The legacy DSP hardware consists of three commercial off-the-shelf (COTS) and 19 custom cards, installed in a single 19" Multibus chassis. This technology will soon be unsupportable.

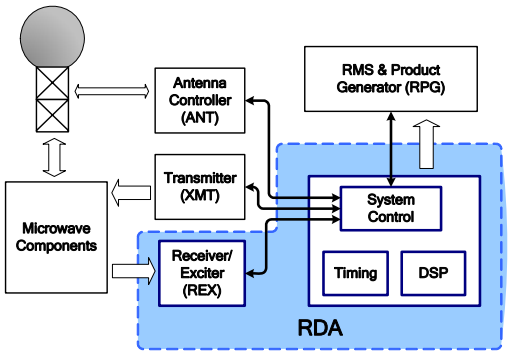

Figure 1. Block diagram of the legacy TDWR.

* This work was sponsored by the Federal Aviation Administration under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the U.S. Government.

Corresponding author address: John Y. N. Cho, MIT Lincoln Laboratory, 244 Wood St., Lexington, MA 02420-9185; e-mail: jync@mit.edu.

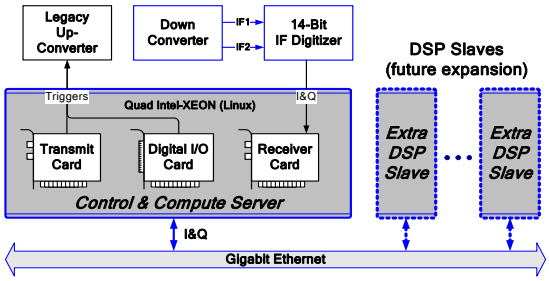

Figure 2. The new RDA can use a cluster of Linux PCs for future enhancements to the digital signal processing algorithms.

A block diagram for the re-hosted RDA hardware is shown in Figure 2. A Quad Intel Xeon (currently 2.7 GHz) processor compute server running Linux performs both the signal processing and system control functions. This computer houses a SIGMET RVP8 system that provides a COTS solution for the digital receiver, digital waveform shaping, and timing functions in three PCI cards, each with several field programmable gate array (FPGA) chips. A combination of interrupt-driven software and FPGA code allows the system to change pulse repetition interval (PRI) sequences and phase coding on a radial-by-radial basis.

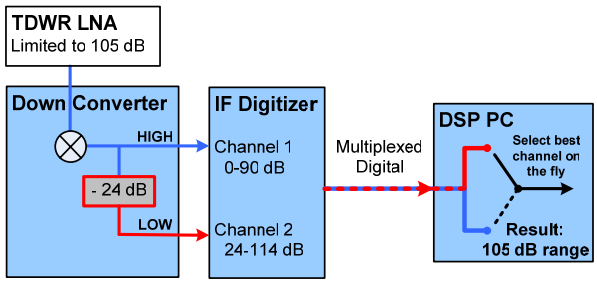

Figure 3. 105 dB of dynamic range is achieved by using two A/D channels.

The TDWR requirement for receiver dynamic range (100 dB) provided a technical challenge. To meet this requirement the legacy receiver used automatic gain control (AGC), which unfortunately degrades system stability and clutter filter performance. Figure 3 illustrates the new RDA receiver design, which uses the SIGMET RVP8 in an extra-wide dynamic range configuration. The RVP8 provides two A/D channels, each with the same dynamic range. The newly designed TDWR down converter feeds these A/D channels with identical

IF signals but offset by 24 dB so that large clutter targets that saturate the high-gain channel can be properly detected by the low-gain channel. A combination of FPGA and C code dynamically measures and then removes the phase difference between the two channels. The RVP8 software switches to the unsaturated stream when necessary at each range gate, pulse by pulse. The resulting dynamic range is now truly instantaneous and is limited only by the upstream TDWR low-noise amplifiers (105 dB). Degradation of data fidelity by the AGC has been eliminated.

Phase stability of the overall system has also been improved by two approaches: 1) with the ability to sample (and cohere to) the phase of the transmitted pulse, and 2) with the replacement of the legacy REX with modern digital pulse forming and receiving systems. The legacy clutter suppression limit due to phase stability was 60.7 dB, whereas we were able to achieve up to 65 dB clutter suppression using a nearby water tower for a target.

A key design feature of the RDA prototype architecture is its scalability, since future algorithm enhancements may exceed the capacity of any single computer system. The legacy signal processing algorithms (requiring less than 250 MFLOPS) can easily be accommodated by one single-processor compute server. The recently developed algorithms discussed in this paper (~1 GFLOP) can also be handled adequately by a single quad-processor system. However, future algorithms under consideration, such as range oversampling, may require an order-of-magnitude increase in processing. The enhanced architecture can use a cluster of COTS compute servers connected via gigabit Ethernet to provide broad scalability to meet the future computational requirements. The in-phase and quadrature (I&Q) data are distributed to a selectable number of DSP slaves using a software standard called Message Passing Interface (MPI). Since all of the software is designed for standards-based interfaces, the final choice of hardware brand and model can be made just before deployment to take advantage of increases in processor speed.

A block diagram of the re-hosted RDA software design is shown in Figure 4 (Elkin and Parker, 2003). System control paths between the legacy RMS, transmitter, antenna, and the new RDA are provided by a message passing process. It provides scan strategies to the antenna, and waveform data and scan commands to the IQ Master (IQM). The IQM sends transmit instructions to the RVP8, accesses the received I&Q data, and serves the I&Q data to the DSP algorithms running on IQ Slave (IQS) processes. Each IQS process operates on a subset of the received range gates and sends moment data to a collector process which reassembles the range slices and outputs the data to the legacy radar product generator (RPG). Insertion of DSP algorithm enhancements is facilitated by the object-oriented design of the IQS that provides an application programmer interface (API).

The engineering prototype of the RDA was developed and tested in fiscal years 2004-5 at FAA's Oklahoma City TDWR Program Support Facility (PSF) site. FAA personnel are currently evaluating and testing an initial configuration that provides hardware replacement of the RDA platform and implements the existing DSP algorithms in software. In 2006 this initial configuration will be fielded at one or two operational TDWR sites. This will provide an opportunity to record I&Q data to facilitate the algorithm development effort. Algorithmic improvements such as those discussed below will be inserted into the platform for evaluation by FAA meteorologists, potentially in parallel with operational deployment of the new RDA design.

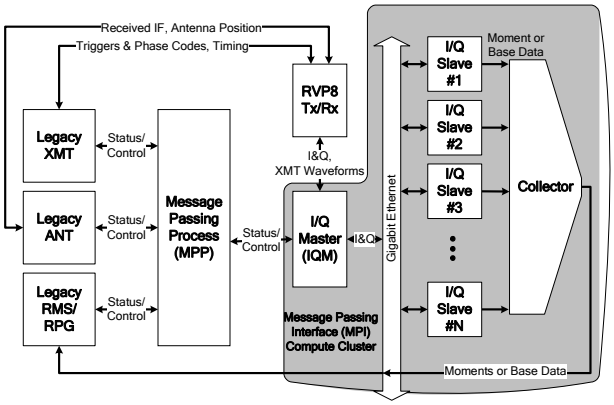

Figure 4. Re-hosted RDA software.

3. RANGE-VELOCITY AMBIGUITY MITIGATION

For TDWRs nationwide, the most serious data quality challenge is range-velocity ambiguity. Because the TDWR operates at C band, this ambiguity is more severe compared to S-band radars like the NEXRAD. Therefore, a very aggressive approach must be taken to combat this problem. On the other hand, the required range for velocity recovery is only 90 km, so the range ambiguity problem is essentially one of first-trip protection rather than the unfolding of multiple trips. The scheme that we devised has been described in past papers, so we will give only a brief explanation here.

3.1 *Algorithm Description*

With the availability of diversity in PRI and pulse-transmission phase, multi-PRI and phase-coding techniques can be applied to this problem. These two approaches have complementary strengths and weaknesses for range-overlay protection (Cho et al., 2003). Multi-PRI signals can be processed to effectively separate different-trip weather even if the overlaid powers are strong or spectrally wide, as long as the overlaid weather does not continuously span a long radial distance. Phase-code processing (e.g., Siggia, 1983; Sachidananda and Zrnić, 1999) works well for trip separation even if the overlaid storm has a long continuous radial range, but breaks down in cases of strong and/or spectrally wide overlays, and also if there are simultaneous overlays from different trips. Therefore, we proposed an adaptive solution where, for low-elevation tilts, information from an initial long-PRI scan

would be used to select multi-PRI or phase-code signal transmission and processing on a radial-by-radial basis in the subsequent scan (Cho, 2003, Cho et al., 2004).

Velocity dealiasing is also necessary to meet the FAA requirement of ±40 m s^{-1} velocity measurement interval, since even at the shortest PRI allowed by the transmitter (518 μs) the unambiguous velocity is only 26 m s^{-1}. For the multi-PRI mode, velocity is dealiased using the unfolded-velocity clustering method (Trunk and Brockett, 1993) on all sets of PRIs that are free of range overlays. For the phase-code processing mode, the dealiasing is carried out across consecutive radial pairs. Again, the clustering method is used so that the PRIs for contiguous dwells need not form a simple integral ratio. After the dealiased velocity field is computed, a 2D (3x3 in range and azimuth) filter is applied to correct instances of false dealiasing (Cho, 2005).

An additional complication to range-overlay protection is the ground clutter filter (GCF). For the multi-PRI mode, applying a GCF coherently across all PRI pulse sets convolves information from different pulses and destroys the independence of range aliasing between PRI sets. Hence, the GCF should only be applied when absolutely necessary. First, apply the adaptive multi-PRI GCF (Cho and Chornoboy, 2005) if clutter signal is present in the long-PRI scan data and all pulses are clean from range-aliasing or the estimated clutter power exceeds the average range-aliased power. Then if the power removed by the GCF is non-negligible, use the filtered time series for further processing; otherwise, use the unfiltered data. Second, if there is non-negligible clutter power present, but the multi-PRI GCF was not applied because of range-overlay contamination, then the means are subtracted from the I&Q components (clean pulses only) of each PRI set separately. In this way, the zero-Doppler power is removed from each PRI set without cross-contamination. Obviously, the clutter suppression is limited in this case, but it is an improvement over no suppression.

Likewise, for the phase-code processing mode, applying a GCF in the first trip severely weakens its ability to filter out the unwanted out-of-trip signal by destroying its phase coherence. Thus, clutter power estimates from the Gaussian model adaptive processing (GMAP) clutter filter (Siggia and Passarelli, 2004) and overlay data from the long-PRI scan are used to decide whether to process the filtered or unfiltered data.

Note that the adaptive transmission and processing scheme will only be applied on low-elevation scans where range-velocity ambiguity is greatest. On high-elevation scans where range aliasing is not a problem (but velocity aliasing is) and ground clutter is minimal, we propose to use dual-PRI staggered transmission. On intermediate elevation scans where the low-level long-PRI scan data are not valid for range overlay determination, only phase-code processing will be applied with alternating-dwell dual-PRI.

Further algorithmic details of adaptive waveform selection, phase-code processing, and data censoring will be presented in a future publication.

3.2 *Example Results*

Although real-time implementation of the adaptive waveform selection algorithm is not yet completed, we recorded I&Q data on the prototype TDWR RDA at the PSF in Oklahoma City for processing later offline. To test the algorithms we collected 0.3° elevation data while cycling through each of the pulse transmission types to be used in the adaptive selection after each 360° scan. The waveforms used were the long-PRI (3066 μs), alternating-dwell dual-PRI (600/900 μs), and three varieties of multiblock-staggered PRI (8x8 [622, 666, 710, 754, 798, 842, 886, 930 μs]; 4x17 [600, 691, 782, 873 μs]; and 4x16 [647, 734, 821, 908 μs]), where multiblock staggered means *m* PRIs x *n* pulses in succession for each PRI. All of these pulse trains employed a pseudorandom phase code on transmission, so that cohering to the first trip renders the other trips incoherent. (Phase-code processing, i.e., spectral filtering of the unwanted trip, is only possible on the alternating-dwell dual-PRI mode.) We also included a single-PRI (600 μs) waveform with constant phase to simulate the legacy mode. The antenna rotation rate was set to 19° s^{-1}, which is the current operational rate used during monitor scans. The TDWR antenna beamwidth is 0.55° with an operating frequency of 5.62 GHz, peak power of 250 kW, and pulse length of 1.1 μs. Range samples are produced at 150-m intervals.

On 14 May 2005 around 0300 UTC, a gust front was propagating across the radar field of view toward the SSE. At the same time, however, there were out-of-trip overlays in the SW to SE sector. With a single PRI and no phase-code transmission/processing (i.e., in a legacy-like mode), the gust front was obscured as it traveled south of the radar (Figure 5, top panel). With the adaptive processing algorithm choosing from phase-code processing and multi-PRI radials, the gust front became distinctly visible in both the reflectivity and velocity fields (the latter shown in Figure 5, bottom panel). In this case the adaptive selection algorithm chose most of the radials to be from the alternating-dwell dual-PRI phase-code processed data, because the out-of-trip weather patches happened to be contiguously long in the radial dimension. The range overlay protection is not perfect, but is a dramatic improvement over the single-PRI case.

To emphasize the extent of obscuration as observed by a downstream user, we also show censored versions of the previous plots in Figure 6. The top panel shows velocity censoring for range overlays as prescribed in the legacy algorithm. Note that a gust front probably would not have been detected by any automated algorithm due to the lack of data in the crucial area. With the results from our adaptive selection scheme, much of the needed information is restored.

Note that our single-PRI, constant-phase mode does not exactly emulate the legacy system. For one, the legacy system uses fixed time-domain IIR clutter filters, whereas we employ an adaptive frequency-domain filter. For another, the legacy system does use information from the long-PRI scan to change the PRI of the subsequent scan. However, since only one PRI can

be selected for all radials, the effect is only to shift the overlays in range and not to eliminate or filter them from the field of view.

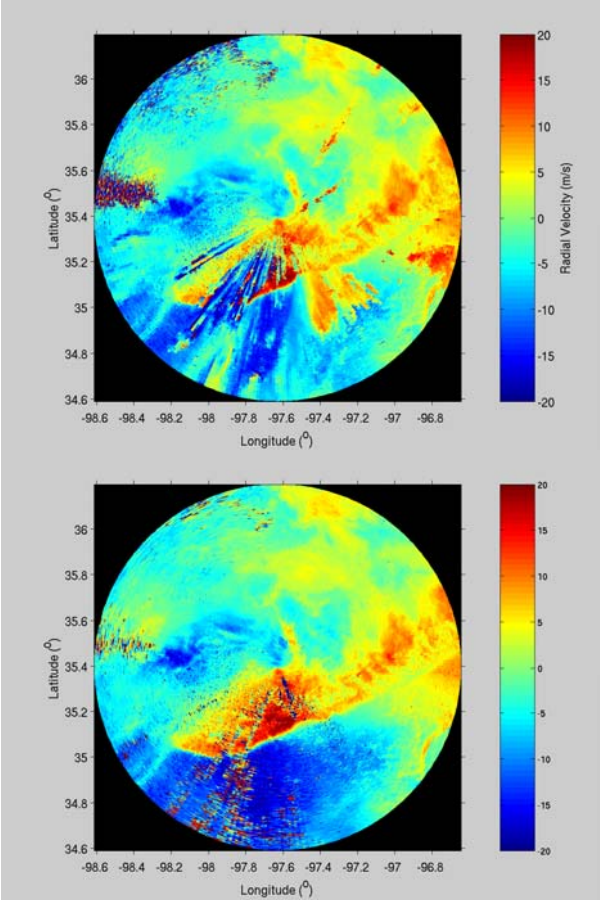

Figure 5. Radial velocity estimated from single-PRI and constant pulse-phase transmission data (top) and adaptive selection algorithm data (bottom).

Velocity dealiasing examples are not shown here, but can be seen elsewhere (Cho, 2005). We have yet to capture a case where severe range and velocity aliasing occur at the same time.

4. SUMMARY

In an effort to enhance supportability, a new RDA system for the TDWR has been developed featuring scalability and open-system architecture. The prototype RDA displays improved performance characteristics relative to the legacy system, especially in instantaneous dynamic range and phase stability. New signal processing algorithms to improve data quality have also been developed, taking advantage of the increased computing power and flexibility in transmitter control. Adaptive selection of multi-PRI and phase-code processing modes results in considerable abatement of the range-velocity ambiguity problem. Work to implement the new algorithms in real time is currently in progress.

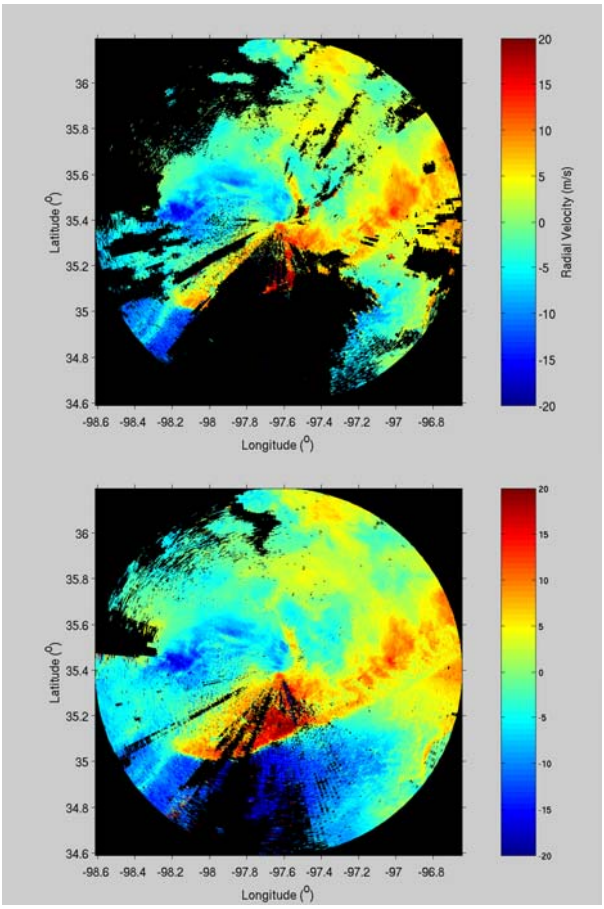

Figure 6. Same as Figure 1 except with data censoring turned on.

5. REFERENCES

Cho, J. Y. N., 2003: Evaluation of TDWR range-ambiguity mitigation techniques. Project Rep. ATC-310, MIT Lincoln Laboratory, Lexington, MA, 47 pp.

——, 2005: Multi-PRI signal processing for the Terminal Doppler Weather Radar. Part II: Range-velocity ambiguity mitigation. *J. Atmos. Oceanic Technol*., in press.

——, and E. S. Chornoboy, 2005: Multi-PRI signal processing for the Terminal Doppler Weather Radar. Part I: Clutter filtering. *J. Atmos. Oceanic Technol*., **22**, 575-582.

——, G. R. Elkin, and N. G. Parker, 2003: Range-velocity ambiguity mitigation schemes for the enhanced Terminal Doppler Weather Radar. Preprints, *31st Conf. on Radar Meteorology*, Seattle, WA, Amer. Meteor. Soc., 463-466.

——, N. G. Parker, and G. R. Elkin, 2004: Improved range-velocity ambiguity mitigation for the Terminal Doppler Weather Radar. Preprints, *11th Conf. on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, Amer. Meteor. Soc., 5.4.

Elkin, G. R, and N. G. Parker, 2003: Terminal Doppler weather radar (TDWR) radar data acquisition (RDA) subsystem upgrade. Preprints, *7th Annual*

High Performance Embedded Computing Workshop, Lexington, MA, MIT Lincoln Laboratory.

Sachidananda, M., and D. S. Zrnić, 1999: Systematic phase codes for resolving range overlaid signals in a Doppler weather radar. *J. Atmos. Oceanic Technol.*, **16**, 1351-1363.

Siggia, A., 1983: Processing phase coded radar signals with adaptive digital filters. Preprints, *21st Int. Conf. on Radar Meteorology*, Edmonton, AB, Canada, Amer. Meteor. Soc., 167-172.

——, and R. E. Passarelli, Jr., 2004: Gaussian model adaptive processing (GMAP) for improved ground clutter cancelation and moment estimation. Preprints, *3rd European Conf. on Radar in Meteorology and Hydrology*, Visby, Sweden, Copernicus Gesellschaft.

Trunk, G., and S. Brockett, 1993: Range and velocity ambiguity reduction. Preprints, *1993 IEEE National Radar Conf.*, 146-149.