Analysis and Automated Target Detection for Passive Bistatic Radar
Fall 2013 EC707
Michael Hirsch & Matthew Kidd

ABSTRACT

Ionospheric fine structure research uses a variety of sensors to study ionospheric events across the (k,ω) space. The heterogenous plasma characteristics (density, temperature, velocity) of the ionosphere are not fully spatially or temporally resolved even by the best radar systems due to finite beamwidth, finite statistical measurement time, and limited access to wavenumber *k* [Lind 2013]. Passive bistatic radar offers opportunities to researchers at little capital expense (starting at $10 for a USB receiver with simple yagi and laptop) by "stealing" existing broadcast signals such that coherent ionospheric events can be detected with compact antennas and receivers.

We present preliminary results on two fronts:

(1) [Matt Kidd] Developing better quantitative metrics for when signals should be used or neglected as part of the target cross-correlation and/or interferometry processes, based on the transmitter signal self-ambiguity function in range-Doppler space.

(2) [Michael Hirsch] Developing a blind target detection system that without ongoing human training can detect events of interest while rarely falsely declaring clutter or interference as valid targets.

INTRODUCTION

MIT Haystack and the University of Washington have been at the forefront of passive radar ionospheric research over the past two decades with the ISIS distributed passive radar instrument. Until recently [Lind 2013], only a few 150kHz segments of the FM broadcast spectrum could be simultaneously captured. Storage and sharing of data has been problematic due to limitations in instrument site storage, processing, and network bandwidth. All three barriers have been pushed down with $130 USB 3.0 4TB hard drives with sustained sequential 100MB/sec read/write speeds, powerful quad-core $500 desktop PCs, and nearly-ubiquitous 50Mbps internet connections. The RF receivers themselves have markedly improved, with recent COTS models allowing 120MHz streaming bandwidth versus the 2MHz streaming bandwidth previously used in the ISIS instrument [F. Lind, private comm. Nov 2013].

We sought to gain insight into the quality of passive radar ambiguity functions that can be obtained from typical FM broadcasts in the United States. Two of the most ubiquitous types of broadcasts are the talk-heavy National Public Radio (NPR) broadcasts and the primarily musical broadcasts of commercial “classic rock” stations. Rock music station aural content consists of mostly musical programming interspersed with commercials and brief periods of broadcaster commentary between songs and advertisements. It was our intuition, and the qualitative experience of our collaborator Frank Lind at the MIT Haystack observatory, that the NPR self-ambiguity functions would be worse than the rock stations due to the relatively long periods of silence inherent in the news and interview format of the NPR station. The FM station self-ambiguity functions, on the other hand, would be of higher and more consistent quality (more thumbtack-like in the range-Doppler space) since the commercial operators have a substantial interest in keeping the pacing of the programming fast and the amount of dead air at a minimum. The aim of this study is to examine how correlated the content of the broadcast is with the self-ambiguity function quality -- for example, how reliably do the pauses in conversations on talk radio cause poor ambiguity function quality compared to a rock station broadcast.

Given the countless terabytes of data collected over the past fifteen years, with an ever increasing data rate collection due to additional sites and broader RF bandwidth, the need for automated target detection is increasingly urgent not only for ISIS, but for other passive radar instruments. The larger the data bandwidth, the larger the need to prune the data early for relevant events due to limited HDD resources. Inexpensive desktop PCs are capable of on-line target detection as we will show in this report, written in straightforward MATLAB or Python script. We use a blind machine vision process that does not require human algorithm “training” or extensive fiddling with parameters. We simply started with values that heuristically made sense, and made only minor if any adjustments to the initial heuristically chosen parameters. We found empirically that some data has highly unstable cross-ambiguity functions with signal to clutter ratio (SCR) varying over the entire dataframe (range-Doppler plot) by several orders of magnitude in a non-stationary way. Such variations would be vexing to a standard machine vision algorithm. A future pathway to better exploiting these data intermixed with sporadic bad data may be on-line qualification of dataframes by measured self-ambiguity of the reference signal. We present machine vision results and MATLAB code at http://heaviside.bu.edu/~mhirsch/isis

BACKGROUND: PASSIVE BISTATIC RADAR

As discussed in [Sahr 1997,Lind 2013], coherent ionospheric returns are readily detectable using passive bistatic radar that exploit FM broadcast signals as the incident radar waveform. Until the past decade or so, extensive detection of aircraft and the ionosphere have been difficult due to RF/data bandwidth and on-line digital signal processing constraints, along with the simple difficulty of transporting gigabytes of data without a high-speed Internet connection [Sahr 1997, Willis]. Some of the early challenges of the high-performance passive radar activated in 1996 by Sahr and Lind included the intense computational load [Morabito] of computing the lag products [Sahr 1997]. Analog NTSC broadcasts contained a strong, low-entropy burst to sync each frame that made detection of overspread targets such as coherent ionospheric returns difficult [Sahr 1997]. Test results from the past decade showed cellular GSM transmission in the 900MHz and 1.8GHz bands to be of insufficient signal strength and RF bandwidth to be useful for ionospheric returns with the passive radars of that time [Willis].

A first step in evaluating the feasibility of passive bistatic radar for detection of a target is considering the power density at the target location due to the FM broadcast transmitter in Eqn.(A-1)[Willis]. Willis notes that the typical highest legal power in the United States is 250kW for FM broadcasts and 1MW for digital TV broadcasts in the UHF channels. Assuming $F_T = 1$, the maximum power density for the Siena College/Dartmouth pairing using the azimuthally omnidirectional WQBJ 50kW ERP transmitter relevant to the 2010-Aug-03 event to be discussed in the Blind Target Detection section of this report is $S_{max} = \frac{50\times10^3}{4\pi(400\times10^3)^2} = 24.9$ nW/m^2 = -76dBW/m^2, where 400km has been selected as the minimum possible range for ionospheric returns between these sites (per Lind, private comm.).

$$S = \frac{P_T G_T}{4\pi R^2{}_T} F^2{}_T \text{ [Watts/m^2]} \qquad \text{(A-1)}$$

where:

P_T is the transmitter output power at the matched antenna terminals

G_T is the antenna gain relative to an isotropic radiator *in the direction of the target* → FM, TV, and cellular transmitters radiation patterns are generally slightly downtilted to focus their energy at the ground as much as possible--passive bistatic radar exploits the sidelobes of these transmitter radiation patterns.

R_T is the distance from transmitter to target

F_T is the pattern propagation factor [Mahafza 2000], which is a catch-all for transmission path losses besides free-space loss accounted for in the $\frac{1}{4\pi R^2}$ term. These losses can be significant in certain bistatic scenarios and include multipath, atmospheric (not very relevant below 1GHz) , diffraction, *et al*.

This transmitter signal when used as a radar waveform has a self-ambiguity function [Woodward, Richards, Mahafza, Willis] defined in Eqn. (A-2).

$$|\chi(\tau, f_d)|^2 = |\int_{-\infty}^{\infty} s(t) s^*(t+\tau) e^{j2\pi f_d t} dt|^2 \qquad \text{(A-2)}$$

Some practical outcomes of Eqn. (A-2) have been described in the ambiguity function analysis section of this report. [Willis] cites a typical FM broadcast bandwidth of 50 kHz, resulting in a $\frac{c}{2B}$ monostatic range resolution of 3 km, while the HDTV bandwidth of 6 MHz results in a $\frac{c}{2B}$ monostatic range resolution of 25 m--adequate for detection of ionospheric turbulence of decameter to kilometer scale. Of course, actual bistatic range resolution computations are a bit more complex than this, but this back-of-envelope estimate builds intuition on the relative quality of transmitter signal candidates to exploit.

The geometric orientation of a target with respect to the transmitter and receiver sites can have a profound impact on the cross-ambiguity function. We do not have the space to discuss the fine points in this report, but let us simply state that for the orientation in Fig. A.1, the cross-ambiguity shape is very nearly of the self-ambiguity [Willis, Tsao]. This does not necessarily mean the geometry of Fig. A-1 is the "best," since perhaps the return signal will have better SNR for another bistatic geometry due to different clutter, illumination, or due to other factors. [Tsao] shows that for geometries approaching and beyond that of Fig. A-2 with a 90 degree bistatic angle, the shape of the cross-ambiguity function is likewise little changed from that of the self-ambiguity function. As the target comes nearer to being between the sites, with a bistatic angle in the 120 degree to 150 degree range as in Fig. A-3, the cross-ambiguity function spreads in range and Doppler and the sidelobes locations themselves move [Tsao,Willis]. In the limiting case of bistatic angle near 180 degrees in Fig. A-4, the cross-ambiguity function becomes very broad [Tsao] and so alternative means of processing target returns may be warranted. This does not mean such a forward-scatter radar is useless--as shown by [Willis], the large forward scattering cross-section of targets can make such radars a useful "trip-wire" detector. For this project, we are more interested in target returns with bistatic angles less than about 135 degrees, where the cross-ambiguity shape is not tremendously changed from the self-ambiguity shape. ISIS can operate with the reference receiver and scattering receiver at one site (e.g. at MIT Haystack), but for greater sensitivity and opportunistic reasons, often ISIS picks up the reference signal *x(t)* at one site and gathers the scattered signal *y(t)* at another site.

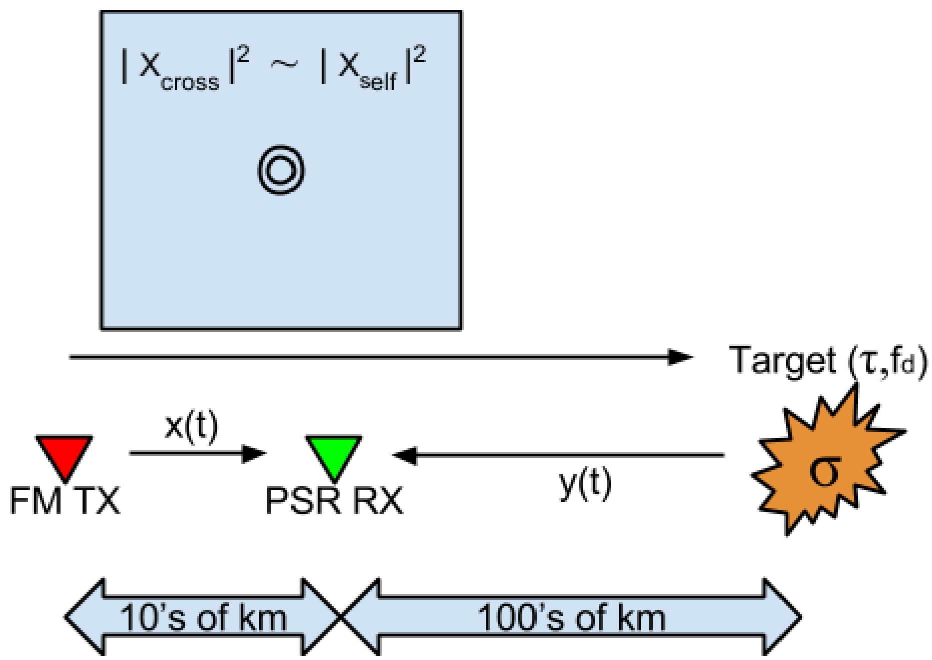

Figure A-1: pseudo-monostatic 0-degree bistatic angle

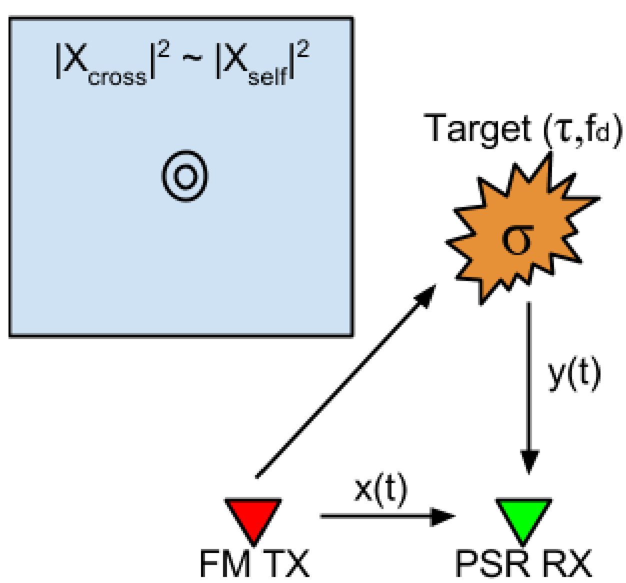

Figure A-2: 90 degree bistatic angle

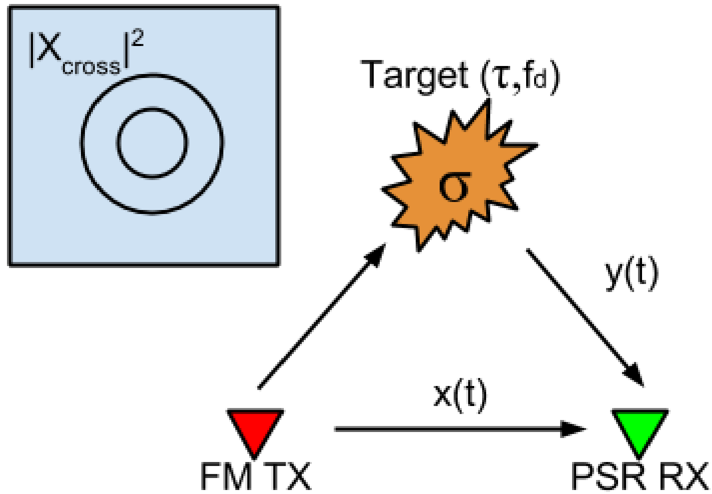

Figure A-3: >120 degree bistatic angle

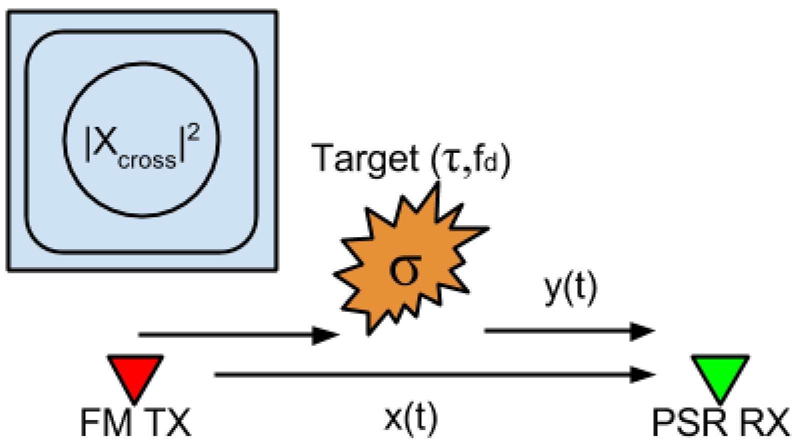

Figure A-4: 180 degree bistatic angle forward scatter geometry

Due to space, scope and IP constraints, we have omitted discussion of the bistatic radar range equation, interference and main transmitter calibration. Certainly these are important issues in system design, particularly with regard to the need for precise time and frequency synchronization of RF sampling [Sahr, Lind 2013]. At this time we forego further such discussion. The publicly available ISIS algorithm block diagram [Lind 2013] is shown in Fig. A-5.

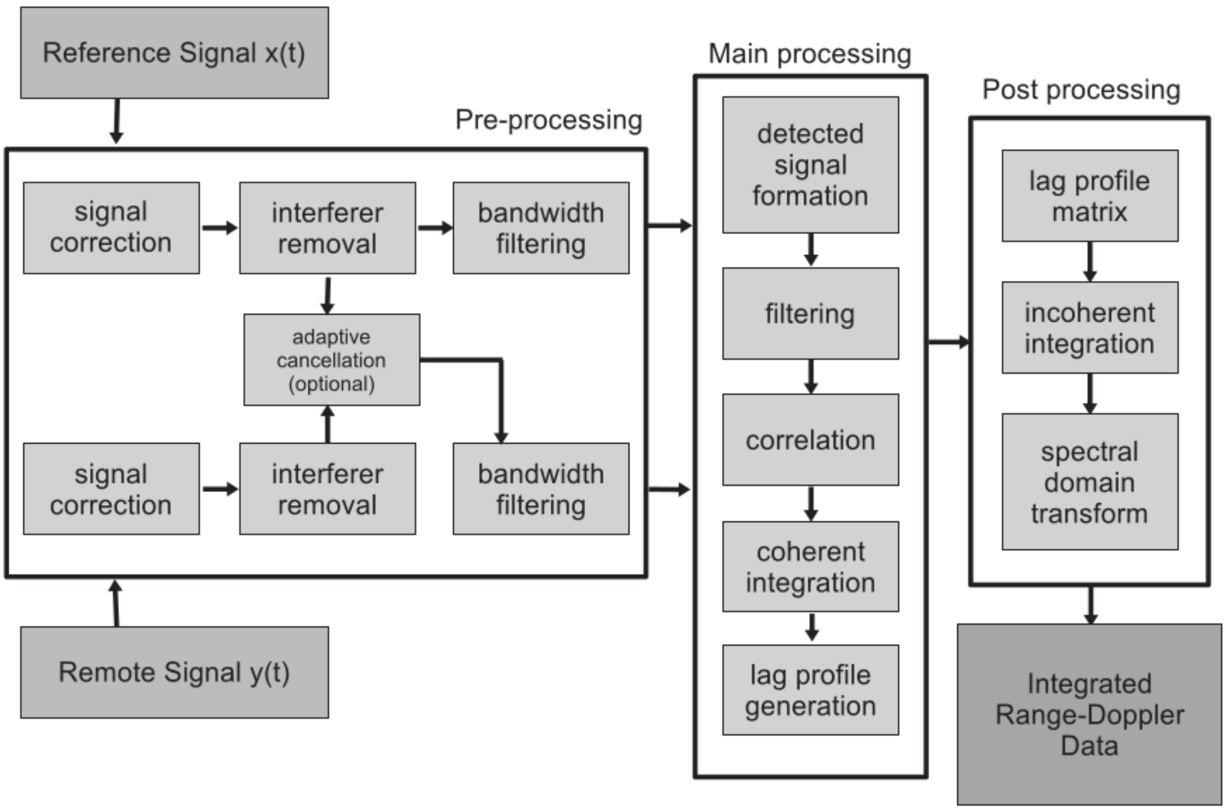

Figure A-5: ISIS signal processing block diagram [Lind 2013]

We conclude the background on passive radar with selected details pertinent to passive radar detection of coherent ionospheric plasma turbulence. [Willis] notes that such turbulence can disrupt waves traveling through the ionosphere up to about 2 GHz, which includes life-critical navigation services such as GPS and aircraft satellite transponders. Better characterization of such turbulent events requires extended data collection. Since we do not know *a priori* when or where such events are occurring, we cannot count on incoherent radar scatter (ISR) sites alone to provide the desired data volume. Since the passive receiver antennas have a relatively broad beamwidth [Lind 2013], the only time limitation is hard drive space (which was a nearly crippling factor in the 1990s [Sahr 1997]) but has become somewhat more tractable with technology advances. Of course, receiver bandwidths and the desire to use multistatic configurations brings the datastream bandwidth up to challenging levels again--necessitating data pruning and curation at an early stage.

A dominant scattering mechanism of these ionospheric turbulences is thought to be Bragg scatter, the effects of which have also manifested in ISR spectrum [Akbari]. Following the development in [Sahr 2007], we define an incident quasi-monochromatic wavevector $\vec{k}_i$ coming from the FM transmitter *x(t)*. Using the assumption [Strømme,Thidé,Sahr 2007] that only longitudinal (along-B) ion-acoustic density waves yield significant scatter, then the scattered wavevector $\vec{k}_s = -\vec{k}_i$ and phonon region wavenumber $\vec{k}_p$ are related by Eqn. (A-3).

$$\vec{k}_i = \vec{k}_p + \vec{k}_s \Leftrightarrow \vec{k}_i = \vec{k}_p - \vec{k}_i \Leftrightarrow 2\vec{k}_i = \vec{k}_p \qquad \text{(A-3)}$$

Observe that $k = \frac{2\pi}{\lambda} \Rightarrow 1/2\lambda_i = \lambda_p$, which is representative of classical Bragg scatter. This means in simple terms that a coherent or incoherent radar will detect only scattering at a wavelength ½ that of the incident waveform *x(t)*. A 100 MHz FM transmitter signal of wavelength 3 meters will yield Bragg scatter from ion-acoustic waves with wavelength 1.5 meters. By inspection we see it is useful to have radars with a wide variety of incident wavenumbers. The ISIS system and upcoming passive radar systems cover the 50MHz to 650MHz range and so should be useful for coherent ionospheric returns at a variety of wavenumbers.

BACKGROUND: FM BROADCAST BAND SIGNALS

As FM broadcast signals are the primary source of radar transmissions for this project, it is useful to briefly discuss their characteristics. For each FM broadcast transmitter, there exists an analytic function *m(t)* comprised of baseband audio (music, voice) covering approximately the frequency range 0..15kHz. The monaural L+R signal of unit amplitude $s_{FM}(t)$ detected by `RawPlayer0.m` is described in Eqn. (B-1) [Lathi]:

$$s_{FM}(t) = \cos(\omega_c t + k_w \theta_m(t)) \qquad \text{(B-1)}$$

where:

$\theta_m(t) = \int_{-\infty}^{t} m(x)dx$ the integral of the modulation *m(t)*

w_c is the FM carrier frequency in radians/sec (e.g. $2\pi \cdot 89.5 \times 10^6$).

t is the independent time variable in seconds.

k_w is the modulation index, representing the ratio of peak deviation to peak modulation frequency--in FM broadcast, $k_w \approx 5$.

x is a dummy variable of integration.

The relatively simple-looking form of $s_{FM}(t)$ in Eqn. (B-1) leads to an RF bandwidth of theoretically infinite extent [Lathi, Lee, Carson]. A formal analysis [Lathi] involving the n-th order Bessel function J_n yields the equivalent form of unit amplitude $s_{FM}(t)$ expressed in Eqn. (B-2).

$$s_{FM}(t) \;=\; \sum_{n=-\infty}^{\infty} J_n(k_w)\cos(\omega_c t + n\omega_m t) \qquad \text{(B-2)}$$

Observe from Eqn. (B-2) that that an infinite number of sidebands exist at frequencies $\omega_c \pm \omega_m, \omega_c \pm 2\omega_m, ..., \omega_c \pm n\omega_m$, and that the spectrum of these sidebands of $s_{FM}(t)$ will decay as $J_n(k_w = 5)$ for the FM broadcast case as shown in Fig. B-1. We repeat here without proof the well-known Carson's Rule [Lathi] for the RF bandwidth B_{FM} containing approximately 98% of the signal energy: $B_{FM} = 2(k_\omega + 1)f_m$, where f_m is the maximum modulation frequency contained in *m(t)*. The practical FM transmitter will likely generate the signal at a frequency much lower than f_c [Lathi], allowing the designer to filter out unnecessary sidebands above and below $n = 1$, since each sideband already contains *m(t)*.

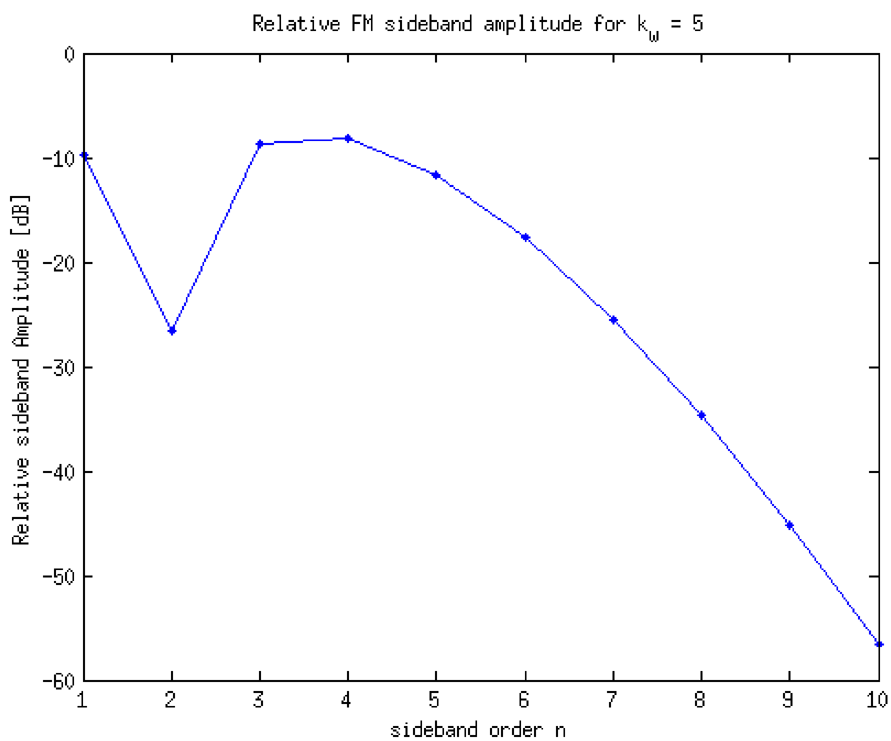

Figure B-1: FM broadcast sidebands relative amplitude

Observe that $s_{FM}(t)$ contains infinitely many copies of *m(t)*. The sideband with the copy of *m(t)* most suitable for demodulation is the lowest order sideband residing at frequency $\omega_c + \omega_m$. The successful demodulation of FM broadcasts in the 88-108MHz band requires:

(1) pre-demodulation low-pass filtering to eliminate the stereo pilot at $f_c + 19$ kHz, the DSB L-R channel from $f_c + 23$ kHz to $f_c + 53$ kHz, and the narrowband FM SCA carriers centered at $f_c + 67$ kHz and $f_c + 92$ kHz.

(2) a near-ideal differentiator yielding $\dot{s}_{FM}(t)$ [Lathi] as in Eqn. (B-3) to recover the modulation *m(t)*.

$$\dot{s}_{FM}(t) = \tfrac{d}{dt}(\cos(\omega_c t \ + \ k_\omega \theta_m(t))) = (\omega_c + k_\omega m(t))\sin(\omega_c t - \pi + k_\omega \theta_m(t)) \qquad \text{(B-3)}$$

By inspection of $\dot{s}_{FM}(t)$ in Eqn. (B-3), observe that the desired *m(t)* is recovered in purely analog receivers by a simple low-pass filter envelope detector, since $f_m \ll f_c$. In actuality, a variety of methods each with their strengths and weaknesses exist for FM demodulation [Lathi]. In the digital domain, several discrete-time solutions have arisen, particularly for embedded systems with limited power/processing resources. The resources needed for low-pass filtering and envelope detection in the digital domain are not trivial for small systems, and so an alternative method has been widely adopted for simple FM demodulation, as implemented in this project. We have been dealing with the real form of unit amplitude $s_{FM}(t) = \Re[\hat{s}_{FM}(t)]$, where

$$\hat{s}_{FM}(t) = s_I(t) \ + js_Q(t) = \ \exp(jk_w\theta_m(t))\exp(j\omega_c t) = \exp(j(\omega_c t + k_\omega \theta_m(t))) \qquad \text{(B-4)}$$

Letting $\phi(t) = \omega_c t + k_\omega \theta_m(t)$ in Eqn. (B-4), we observe that *m(t)* is also accessible by the derivative of the complex phase $\phi(t) = \tan^{-1}\frac{s_Q(t)}{s_I(t)}$. From elementary calculus and algebraic manipulation detailed in [Lyons], we obtain an algorithm for this derivative in Eqn. (B-5).

$$\frac{d}{dt}\phi(t) = \frac{s_I(t)\frac{d}{dt}s_Q(t) - s_Q(t)\frac{d}{dt}s_I(t)}{s_I^2(t) + s_Q^2(t)} \qquad \text{(B-5)}$$

The demodulator we implemented as a first pass in MATLAB used a central difference approximation to the derivative operation. Rather than implementing M-tap FIR filters for the N-point central difference stencil, we simply took the 3-point stencil central difference of the arctangent of Q/I. This gave audibly less distortion than both the 2-point stencil forward difference and the 3-point central difference implementation of the Lyons algorithm. Specifically we entered into MATLAB:

```
m = central_diff( unwrap( atan2(Q,I) ), Ts);
```

where:

`central_diff()` is a 3-point stencil central difference algorithm written by [Canfield].

`unwrap()` is the built-in MATLAB function allowing angles to continue beyond $[-\pi, \pi]$.

`atan2()` is the built-in MATLAB function computing four-quadrant arctangent.

`Ts` is the uniform sampling interval $\rightarrow dt$

$Q = s_Q(t)$ and $I = s_I(t)$

Note that arctangent is a computationally expensive operation--a DSP designer may use industry methods for creating efficient stencils suitable for sufficiently accurate for the digital difference operation.

We invoke the fundamental theorem of calculus in Eqn. (B-6), $\frac{d}{dt}(\omega_c t + k_w \int_{-\infty}^{t} m(x)dx) = \omega_c + k_\omega m(t)$

(B-6)

where by inspection ω_c is a nuisance DC offset and k_w is a constant scalar, both of which are trivially removed from *m(t)*. The demodulated monaural content may audibly add to the reader's intuition when viewing plots of the self-ambiguity discussed in the ambiguity function analysis section.

DESCRIPTION OF METHODOLOGY: AMBIGUITY FUNCTION ANALYSIS

The signals of FM stations 106.7 WIZN and 89.5 WVPR were measured using an ISIS array receiver at 150 kHz RF bandwidth for a period of one hour each by Frank Lind of MIT Haystack Laboratory. Frank Lind processed the data to range and doppler plots by performing the ambiguity calculation as has been described in detail by Sahr *et al* [1997]. The ambiguity functions were averaged over two second intervals of FM broadcast, this being the minimum amount of time required to generate adequate-quality ambiguity functions. Both the ambiguity matrices and the I and Q data were saved to proprietary binary files, and Lind provided us with the Python code to convert these into more convenient HDF5 MATLAB files. Unfortunately, there was an error in the data collection that resulted in much of the NPR data unusable; while all 60 minutes of the rock broadcast were successfully processed into ambiguity functions and I and Q data, only 17 minutes of the NPR data ended up being usable. This limited data, however, seems to have been enough to do some meaningful analysis.

Two metrics were applied to these data, the first being the instantaneous bandwidth of the I and Q data, shown in Eqn. (C-1). Eqn. (C-1) provides an estimation of the bandwidth of the FM signal at any given sample; this approximation has been discussed by Barnes [1992]. Lind uses this metric to determine whether or not to exclude a given sample of the FM broadcast from the ambiguity analysis -- if the instantaneous bandwidth is below a certain threshold, the sample is excluded from the ambiguity analysis. In certain circumstances this practice allows for otherwise unusable broadcasts to be used, although the number of samples excluded can reach 50% in order to generate a high enough quality ambiguity function [Lind 2013].

$$B_i(n) = \frac{1}{8\pi}\left|\ln\frac{I^2(n+1) + Q^2(n+1)}{I^2(n-1) + Q^2(n-1)}\right| \qquad \text{(C-1)}$$

The second metric we used was the the peak-to-side-lobe ratio (PSLR) [Mahafza]. This is a standard measure of ambiguity function quality, calculated by simply taking the ratio of the magnitude of the most intense peak of the ambiguity function to the second most intense side lobe peak expressed in dB.

Finally, the I and Q data were demodulated to monophonic audio in MATLAB using the function `RawPlayer0.m`, and this audio was synchronized with the ambiguity functions and updating plots of the PSLR and instantaneous bandwidth of the broadcast signal. A figure with three subplots was generated using MATLAB; the top subplot is a graph of PSLR and instantaneous bandwidth averaged over two seconds and 0.5 seconds, respectively. The bottom left graph is a graph of an estimate of the power spectral density of a half second of the audio signal, plotted from 0 to 20 kHz generated using the MATLAB function periodogram. The bottom right image is the range-doppler ambiguity of two seconds of the radio broadcast. These figures were generated in half-second increments and synchronized with the demodulated audio, and MATLAB was used to interleave the audio and video frames, creating an AVI video. A typical frame of the video is shown in Fig. C1, and selected videos can be found in the supplementary materials of this paper.

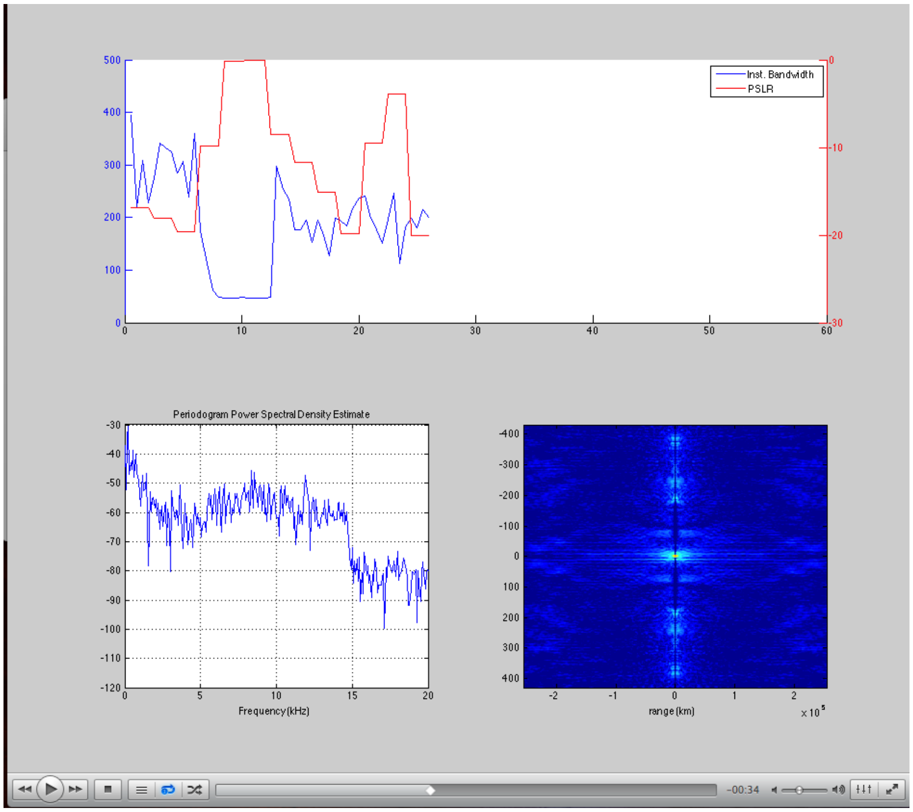

Figure C1: typical video frame

RESULTS: AMBIGUITY FUNCTION ANALYSIS

While much of our analysis was largely qualitative in nature -- matching up features of the audio of the FM broadcast with changes in the ambiguity functions through watching the videos we produced, we were able to do some quantitative analysis of these data. But first we will examine some of the trends and more interesting phenomenological observations we were able to glean from the qualitative video analysis.

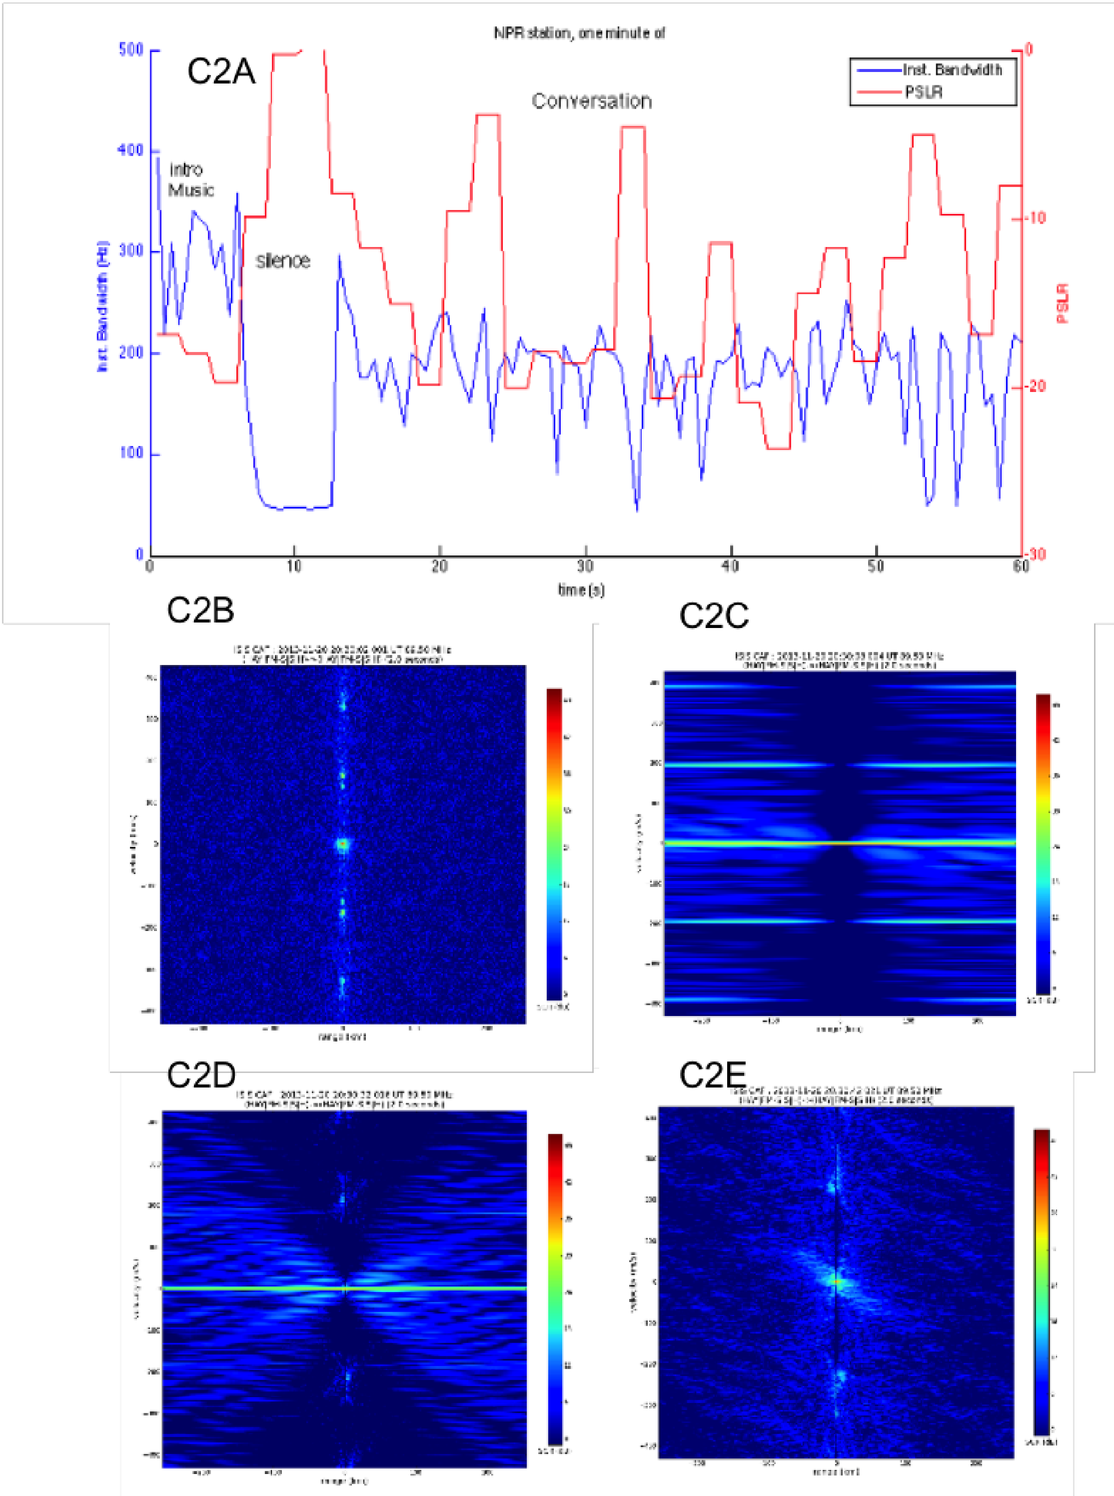

Figure C2, characteristic ambiguity functions and the PSLR and instantaneous bandwidth of a one-minute excerpt of an NPR broadcast

Figure C2a shows the PSLR and instantaneous bandwidth of the first minute of the NPR program. This figure nicely encapsulates the relationship of the content of the broadcast to quality of the ambiguity functions. The first seven seconds of the broadcast are the intro music of the program, followed by 4-5 seconds of dead silence, followed by 45 seconds of a back-and-forth interview between the host and the guest on the program. Characteristic ambiguity functions for the intro music and silence are shown in figure C2b and C2c, respectively. The introductory music produced a good-quality ambiguity function, with the PSLR approaching -20 dB, while the silence produced an awful ambiguity function, with PSLR approaching 0 dB. Once the host and guest began their conversation, the ambiguity function fluctuated between good and poor quality ambiguity functions. Clearly, the long period of silence towards the beginning of the broadcast correlated with the poor ambiguity function produced during that period. The behavior of the ambiguity function PSLR and the instantaneous bandwidth of the broadcast during the conversation phase is typical of conversations throughout the NPR broadcast -- periods of good ambiguity performance are interspersed with poor ambiguity functions, corresponding to pauses in the conversations. Many of these pauses can be seen in the conversation phase of figure 2, where the half-second averages of instantaneous bandwidth drop down near the level of the silence in the first portion of the audio. Characteristic images of bad-quality and good-quality ambiguity functions are shown in figure C2d and C2e, respectively. Note that the good-quality ambiguity function in C2e is not as good as the ambiguity function of the intro music phase, and the bad-quality ambiguity function in C2d is not as bad as the ambiguity function of the silence phase.

Figures C3 and C4 show a plot of instantaneous bandwidth and PSLR for the rock and

NPR station, respectively. Note in figure C3 the minimums of the instantaneous bandwidth at approximately 65 seconds and 335 seconds. These decreases correspond to brief intervals between songs where the disc jockey (DJ) is speaking. Note that the instantaneous bandwidth during these speaking intervals are still quite high compared to the NPR broadcast in figure C4, and the PSLR remains relatively good through these intervals. This is typical of the the rock broadcast -- the instantaneous bandwidth of the rock broadcast rarely drops to the low levels (~200 Hz) that is typically seen throughout the NPR broadcast. The sudden increase in instantaneous bandwidth around 280 seconds in figure C4 correspond to a clip of the guest's television program being played. This clip was much "busier" than the preceding and succeeding conversation -- the dialogue is fast paced, there is constant low background noise, etc. The PSLR of the ambiguity functions correspondingly increases in this region of the NPR broadcast.

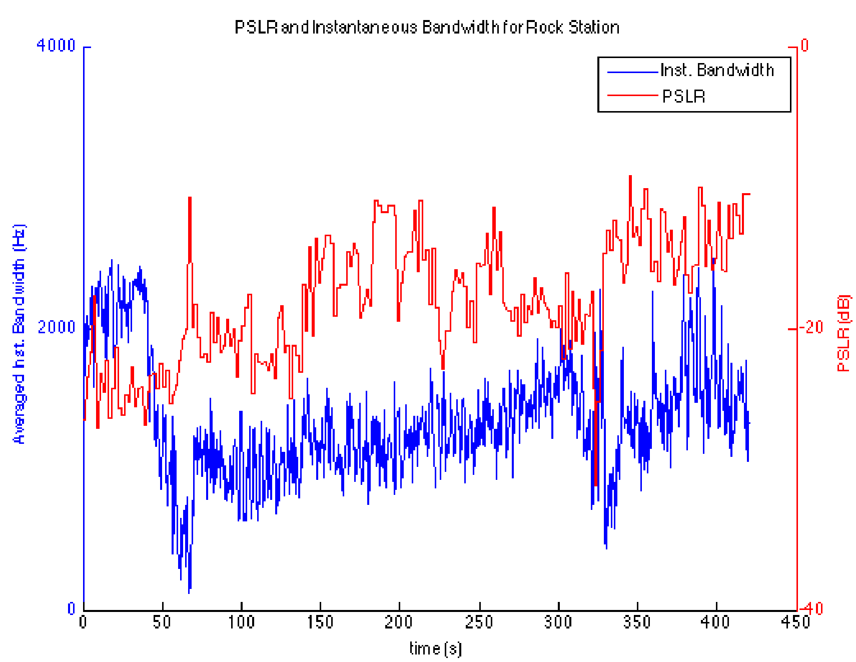

Figure C3: PSLR and average instantaneous bandwidth of a typical stretch of a rock broadcast

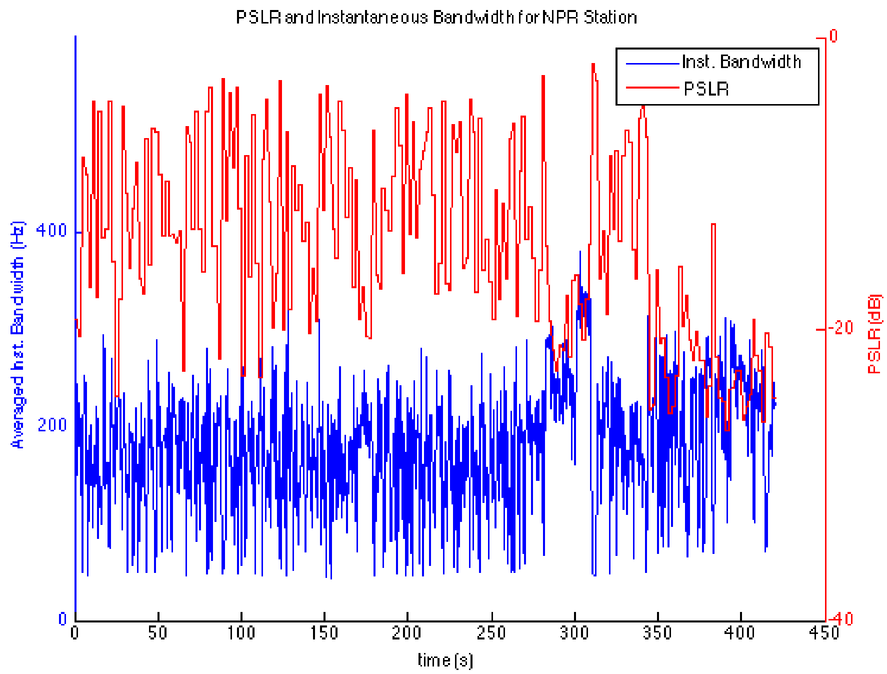

Figure C4: PSLR and average instantaneous bandwidth of a typical stretch of NPR broadcast

Table C1 summarizes some basic statistics of the two broadcasts and their ambiguity performance. These statistics confirm both our intuition and the conclusions drawn from the more qualitative discussion above. The average PSLR is better for the rock broadcast, and perhaps even more importantly, the PSLR variance for the NPR broadcast is more than double the rock broadcast. So, not only does is the NPR station have a higher average PSLR, there is more variability as well. This manifests in the number of 2-second ambiguity functions with PSLR above a certain threshold -- practically no ambiguity functions for the rock broadcast have a PSLR above -10 dB, while more than a quarter of the ambiguity functions of the NPR broadcast are above -10 dB. This is likely explained by the large number of pauses in conversation in the NPR broadcast, while there are no such considered silences in a typical rock broadcast.

Table C1

	NPR	Rock
Average PSLR	-14 dB	-18 dB
PSLR variance	44	19
Average Inst. Bandwidth	179 Hz	1350 Hz
PSLR > -5 dB	10%	0.02%
PSLR >-10 dB	31%	0.5%
PSLR > -10 dB	51%	27%

It is curious, however, that the instantaneous bandwidth of the rock station never reaches the low levels of the NPR station, even during DJ interludes -- and that the NPR instantaneous bandwidth never reaches the levels of the rock station, even during musical segments. It is difficult to say with confidence why this is, but one might speculate that this could be due to the sorts dynamic range compression that a commercial FM station might employ in order to increase the perceived loudness of the broadcast while still remaining in the legally defined loudness limits. On the other hand, an NPR station has little need for this sort of perceived loudness gamesmanship when trying to attract listeners given the more contemplative style of the broadcast.

Finally, figure C5 presents a scatter plot of the averaged instantaneous bandwidth and

PSLR of the NPR station and the rock station. These data sets are mostly distinct in instantaneous bandwidth, and the much higher likelihood of the NPR having a poor ambiguity function is clearly illustrated by this chart. But perhaps most telling are the correlations of these two data sets. First consider the NPR data -- here, the correlation between PSLR and instantaneous bandwidth is negative. This is exactly what one would expect, that the ambiguity functions would be increasingly likely to deteriorate as the randomness of the signal decreased. A linear fit of the NPR data shows a slope of -0.041 dB/Hz and an intercept of -6.8 dB. On the other hand, the rock station's ambiguity function shows no such negative correlation, rather, the linear fit to these data have a positive slope of 0.0014 dB/Hz. The lack of correlation between instantaneous bandwidth and PSLR for the rock station suggest that above a certain threshold the instantaneous bandwidth of the broadcast is no longer an important factor, and that some other factor besides instantaneous bandwidth can determine the quality of the ambiguity function.

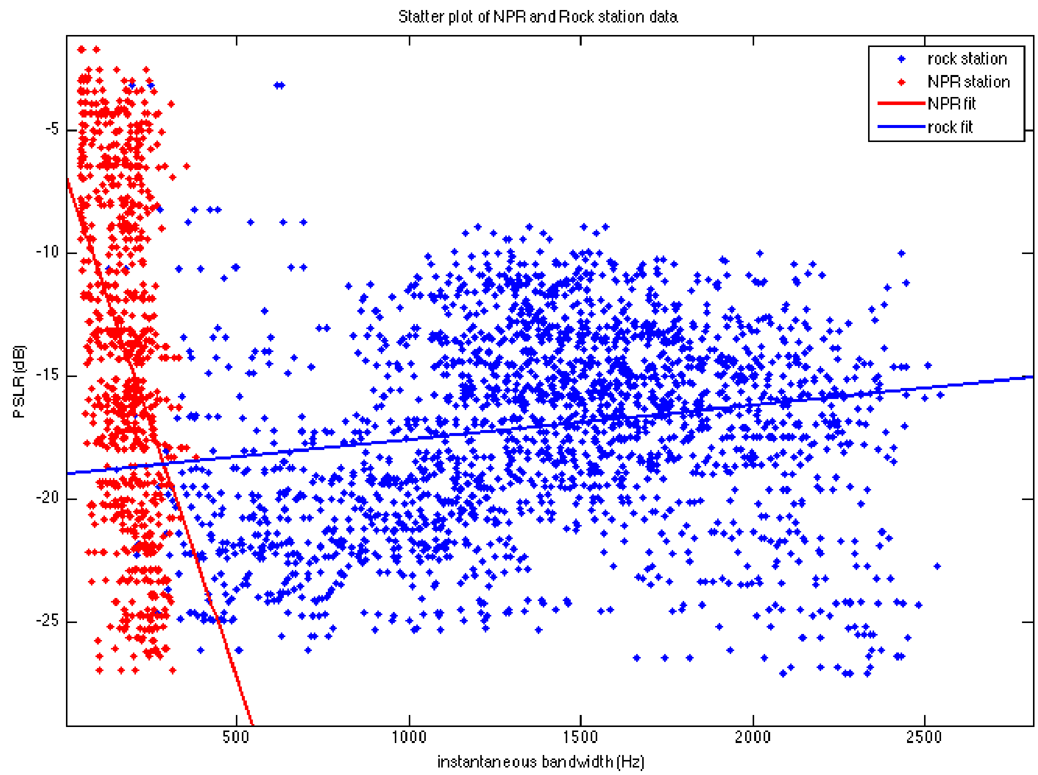

Figure C5, scatter plot of NPR and rock station data, PSLR versus instantaneous bandwidth.

If the reader is interested in watching representative samples of the broadcast described above, some of the audio/ambiguity function videos have been made available in the supplementary materials of this work.

CONCLUSIONS: AMBIGUITY FUNCTION ANALYSIS

While previous to this work it was known that low-instantaneous bandwidth samples can be discarded in order to provide improved ambiguity function, these results provide a framework for a more detailed and considered analysis of how the content of the audio broadcast affects the ambiguity performance, as well as a more detailed understanding of the general ambiguity performance that can be expected from each of these types of broadcast. While researchers in this field already had an accurate heuristic understanding that music channels are more desirable for passive radar work than a talk-heavy format, we have implemented tools that allow for visualization and characterization of these sorts of data. This work aimed to provide more context to the heuristic understanding of the relationship between different radio broadcasting formats, and in that goal we were successful.

In the future, similar analysis could be undertaken of other formats, such as country-western FM stations, AM talk radio and FM talk radio. Similarly, digital television signals could be analyzed to see if different broadcast formats have characteristic ambiguity performance. Additionally, other measures of ambiguity function quality might be better to use -- the author used PSLR in part because of the ease of implementing this standard measure in the compressed timeframe of the project. Integrated side lobe or an estimate of the range and doppler resolution for the signal might be a more appropriate metrics.

DISCUSSION OF IONOSPHERIC EVENTS: 2010 Aug 03

This particular event was sensed by two receive sites each using a distinct transmitter. The Siena College → Dartmouth configuration used WQBJ 103.5FM, a 50kW ERP transmitter in Cobleskill, NY at 116 m elevation. WQBJ is 66.5 km from the Siena College reference receiver at an azimuth of 295 degrees. WQBJ is 196 km from the Dartmouth scatter receiver at an azimuth of 65 degrees. Given that the ionospheric turbulences of interest are expected to exist at 100 km altitude or more, a geometric argument is made that the minimum possible range at which these turbulence returns will occur is approximately 400 km [Lind, private comm. Dec 2013] for the SNC→DART configuration that we examine for the 2010 Aug 03 event.

A list of known events along with the current availability of data is given in Table E-1. The meteor trail event of 2011 July 20 is visible for exactly one 1 second integration time dataframe. This type of event is detectable by simple machine vision techniques using histogram thresholds (e.g. Otsu histogram thresholding and ROI qualification) but we sought to implement more challenging detections requiring a process with "memory" of certain characteristics previous frames that inform decisions on the present frame. We focused on the 2010 Aug 03 event since we had cross-ambiguity data from two site pairs as well as having the raw data immediately available for download.

Event Type	Event Date	Data Ready to Download	Range Location [km]	Doppler Extent [m/s]
Aircraft	2011-Jul-18	No	10 .. 200	-100 .. 100
Meteor trail	2011-Jul-20	Yes	< 10 each	-25 .. 75
Ionosphere	2006-Dec-14	No	700 .. 800	-500 .. 0
Ionosphere (SNC→DART)	2010-Aug-03	Yes	700 .. 900	-800 .. -100
Ionosphere (HAY→ SNC)	2010-Aug-03	Yes	600 .. 850	0 .. 300
Ionosphere (HAY→ HAY)	2012-Mar-09	Yes	500 .. 900	-400 .. -50

Table E-1: ISIS Selected Event List

The basic geometry of the 2010 Aug 03 event is shown in Fig. E-1, with the background color elevation data coming from the ⅓ degree SRTM elevation dataset. Observe the dashed lines to the cartoon blast representations of the ionospheric return. It is possible to postulate that the return came from northwest of DART rather than south of DART since the DART antenna is directed in a northerly direction. Certainly it is possible for a return to leak through a sidelobe of the DART antenna; to get an unambiguous result one would have to invoke additional transmitter and/or receiver sites to establish where on the isocontour the true return lies. Additionally, we have drawn an arrow directly between WQBJ and DART to represent the self-interference that is virtually always a significant factor in passive broadcast radars of this type, but which for space and scope constraints we have omitted from this report besides showing them in Fig. D-1.

The all-important cross-ambiguity function is computed from the digitally sampled $x(t)$ and $y(t)$ as stated in Eqn. E-1 [Lind 2013]. This equation has been deeply analyzed in [Sahr

1996], and given that many of the actual implementation details are proprietary and unknown to us, we will not be able to discuss Eqn. (E-1) in detail. The consequence of Eqn. (E-1) can simply be stated as the cross-ambiguity function is the self-ambiguity function shifted in range and Doppler to the location of a simple point scatterer, and scaled by the RCS of the scatter (along with propagation losses). For targets such as the ionospheric turbulence, we will see a range of Doppler as per discussion in [Sahr 1996, Lind 2013]. We regret not having the capability to describe this process further in this report.

$$\chi[\tfrac{r}{2},\tau] = \sum_{t} y[t]x^*[t-r]y^*[t-\tau]x[t-r-\tau] \qquad \text{(E-1)}$$

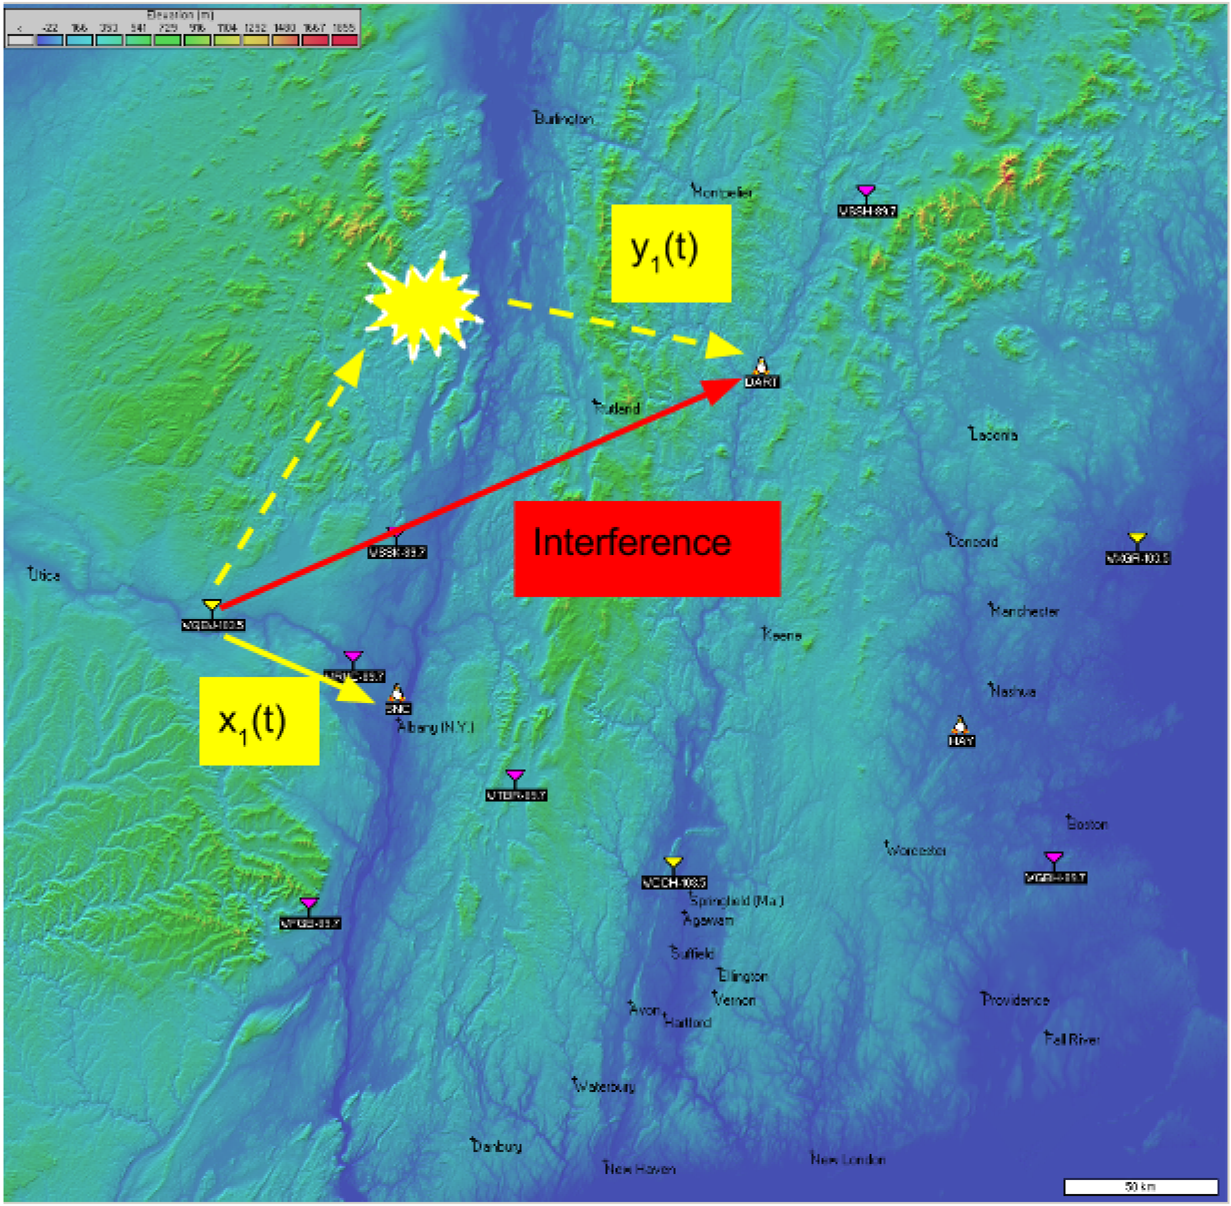

Figure E-1: Overview of 2010 Aug 03 event bistatic geometry

DESCRIPTION OF METHODOLOGY: BLIND TARGET DETECTION

We will show by example that machine vision techniques typically developed for 2D CCTV camera images may be applied to 2D range-Doppler maps that have no relation to traditional camera data. For convenience and consistency, we will refer to the 2D range-Doppler map from a single incoherent integration interval (e.g. 2, 5, or 10 seconds) as a "dataframe." The particular machine vision process focused on by this project is *segmentation*. The overall process is shown in highly simplified form in Fig. F-1. A more detailed view of the preliminary machine vision algorithm *for each pixel* is presented in Fig. F-2. Fig. F-2 does not show the preprocessing step of 2D Wiener filtering that smoothes out clutter in the SCR dataframe.

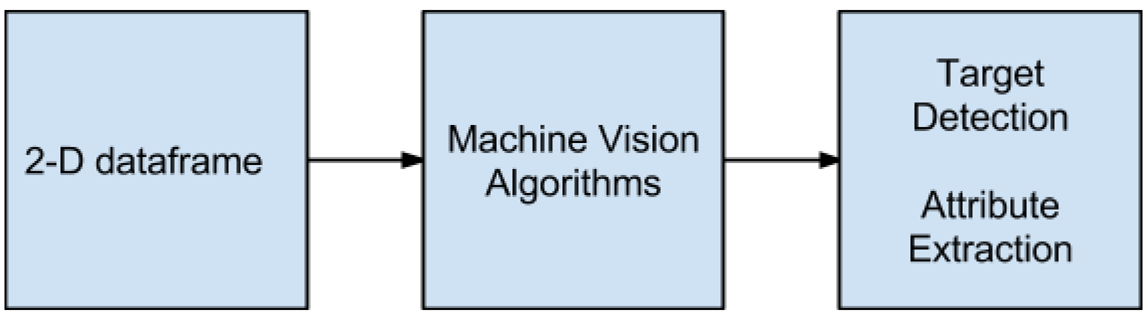

Figure F-1: Overall passive radar data flow algorithm

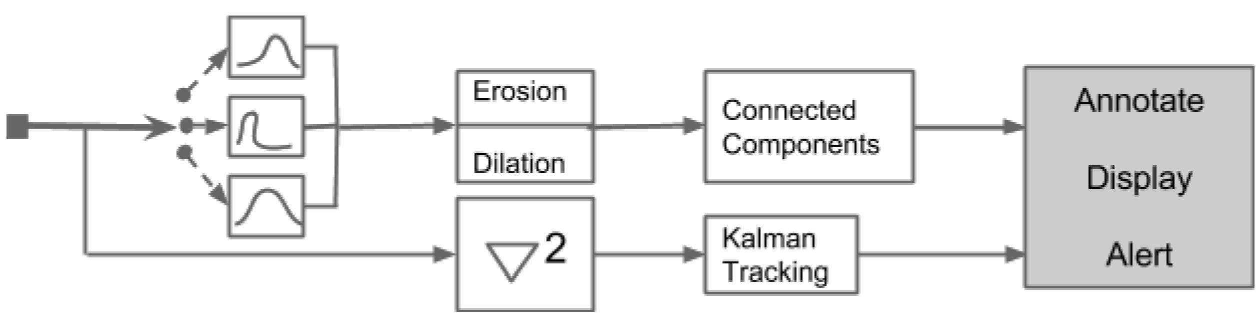

Figure F-2: Detail of candidate machine vision algorithm

The basic goal of segmentation is to divide a dataframe into two or more regions [Gonzalez], e.g. target and clutter or foreground and background. At least one of these dataframe regions would typically be denoted “background,” that is, comprising regions not of further interest, and at least one other dataframe region would be denoted “foreground,” that is, comprising regions of interest for further processing, classification and tracking. Segmentation is rarely a single step process, particularly for dataframes consisting of low SNR data.

GAUSSIAN MIXTURE MODEL:

Originally reported by Stauffer and Grimson [1999], the Gaussian Mixture Method (GMM) goal is to eliminate a majority of the background pixels. Typically GMM implementations experience a few false positives, particularly for strongly non-stationary noise, which manifest as isolated pixels falsely declared as foreground. Again, we stress that GMM is a per-pixel process and does not incorporate information from adjacent pixels in the basic form implemented here. A good example of a raw GMM output exhibiting these false positives on the edges of the tree shadows, while mostly correctly detecting human activity is seen at:
http://www.youtube.com/watch?v=rCTqOYFSEmA
Further processing is almost always needed to obtain a more complete dataframe segmentation. The block diagram showing the decision process between N=3 Gaussian distribution is shown in Fig. F-3. An example frame of MATLAB output is shown in Fig. F-4. Notice in Fig. F-4 that there are a substantial number of pixels falsely declared as foreground (the white pixels in the binary image), despite the 2D Wiener filter pre-processing. The subsequent steps outlined in Fig. F-2 will work to eliminate these false single-pixel declarations.

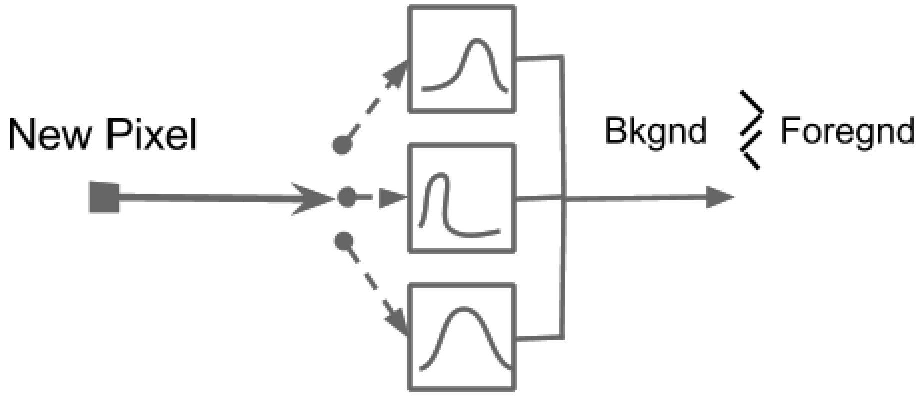

Figure F-3: Gaussian Mixture Model with background/foreground decision output

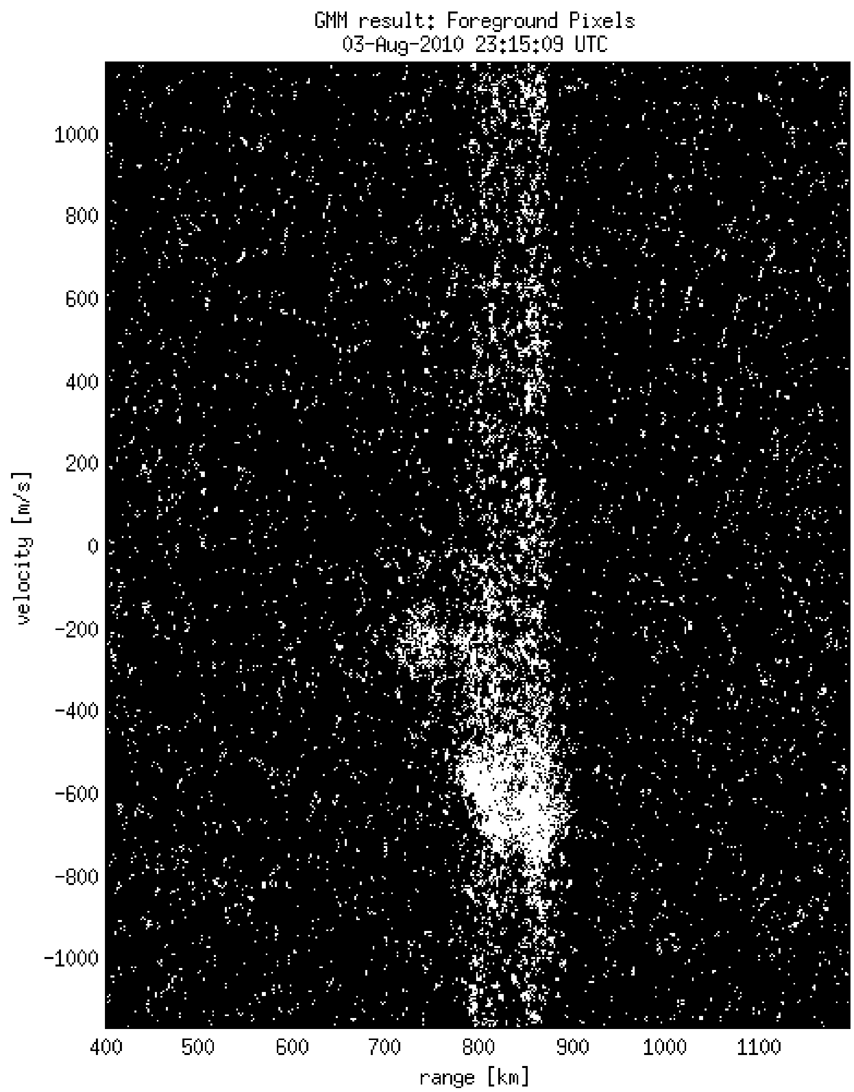

Figure F-4: Example GMM output for actual ionospheric return

MORPHOLOGICAL OPERATIONS:

Morphological operations are machine vision operations using set theory to modify a dataframe on a pixel-by-pixel basis. In this project we work exclusively with morphological erosion and dilation. Morphological *erosion* is a set process in which a structuring element, in this case chosen to be a disk 3 pixels in diameter, is passed over each and every binary pixel region in the image. If such a region cannot completely contain the structuring element, that region is eliminated from the binary image. If a region can contain the structuring element, then as the structuring element is slid around inside the region, the pixels "touched" by the center pixel of the structuring element are preserved. Fig. F-5 shows a cartoon version with the rectangular structuring element shown in orange, and the erosion output as the one-pixel wide horizontal red line. Note, Fig. F-5 cartoon is not representative of the actual structuring element, but was chosen for clarity of the cartoon example. The erosion of actual dataframes is shown in Fig. F-6 from MATLAB. Observe in Fig. F-6 that virtually all false positive foreground declarations have been eliminated. However, the desired convex hull of pixels representing the actual ion-acoustic turbulence return have been eroded down to the point that connected components analysis will miss the associations of nearby pixels and declare a false negative--that no ionospheric turbulence existed here. We must reassociate the pixel regions by performing the next step in the Fig. F-2 process: morphological dilation.

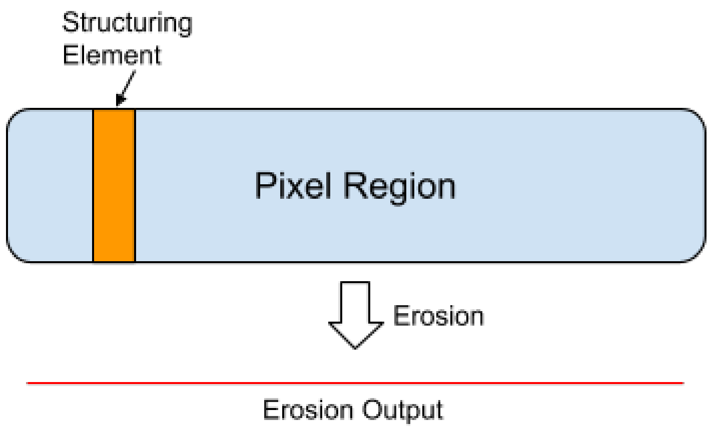

Figure F-5: Morphological erosion example

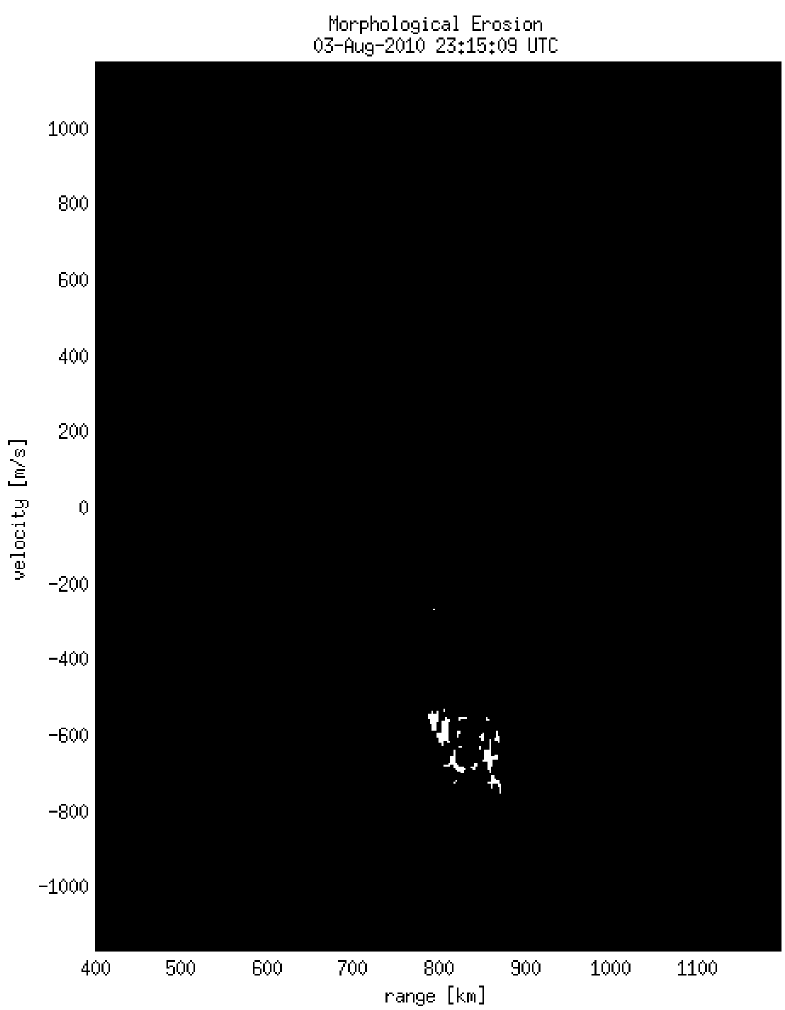

Figure F-6: Example morphological erosion for actual ionospheric return

Morphological *dilation* is a set process in which a structuring element is passed over pixel regions. Imagine a peg in the center pixel of the structuring element--the result of dilation is that any place "touched" by any pixel of the structuring element is declared binary 1. A cartoon example of morphological dilation is shown in Fig. F-7. Note that even if the red line in Fig. F-7 had a break up to one-half the diameter of the structuring element, the continuous green form seen at the bottom of Fig. F-7 would result. This critical fact is exploited to join associated regions of ionospheric turbulence in the processing of real passive radar data. The morphological dilation of actual data in MATLAB using a disk structuring element of diameter 5 pixels is shown in Fig. F-8. Observe how associated pixels regions of the ionosphericic returns have been rejoined. Now the data looks ready for connected component blob analysis, the next step in Fig. F-2.

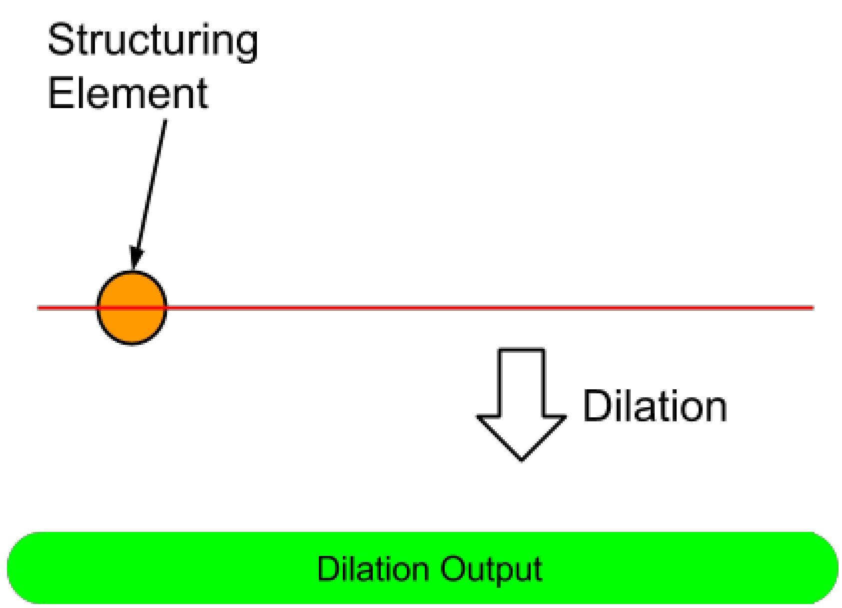

Figure F-7: Morphological dilation example

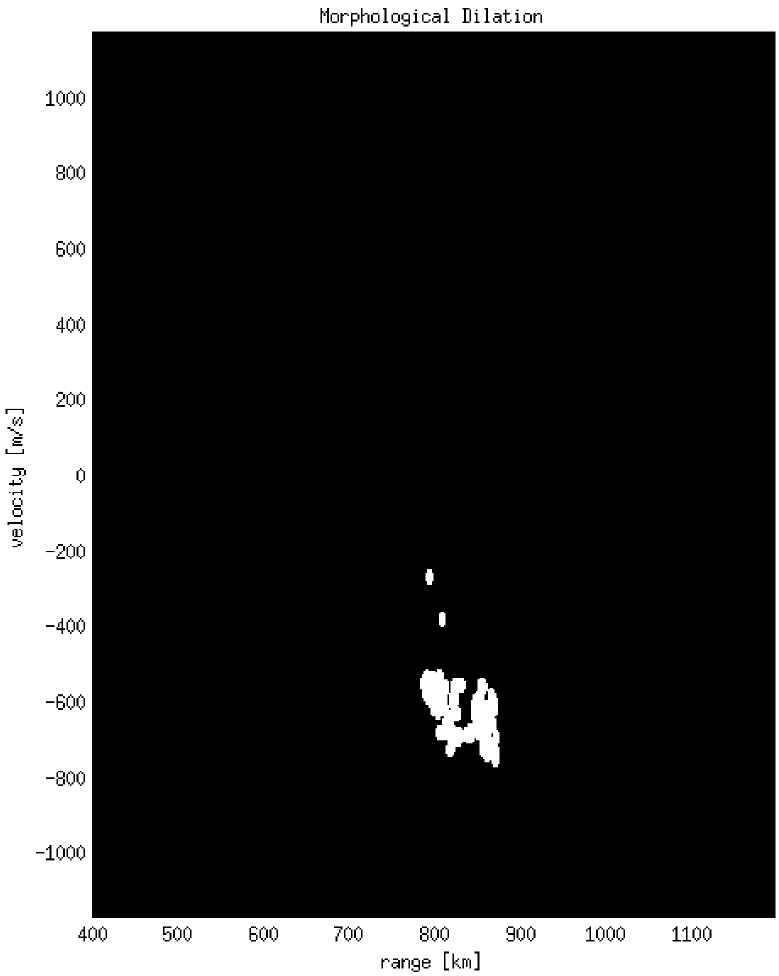

Figure F-8: Example morphological dilation output for actual ionospheric returns

CONNECTED COMPONENT and BLOB ANALYSIS:

To make a final declaration on ionospheric target candidates, we consider whether a region of sufficient associated pixel extent is observed. This might seem to exclude small-scale events--which for now it will--until we gain confidence that not too many false positives are generated over a larger trial set of data. At such a point, we could then consider the time dimension via Kalman tracking or further ROI qualification to better positively classify small-scale ionospheric targets. We declare as connected any region of contiguous 8-connected pixels. 8-connected neighbors mean any pixel touching a face or corner of another pixel as depicted in Fig. F-8.

Blob analysis uses the connected component regions of pixels and computes the size of the oriented minimum bounding box (rectangle) containing the convex hull of a region of associated pixels representing the declared foreground (target) pixel [Gonzalez]. The cartoon depiction of the convex hull (outlined in purple) of connected component pixels is enclosed by the minimum bounding box (outlined in green) as the output of the blob analysis in Fig. F-9. Observe how the upper left group of connected pixels has a convex hull, but does pass blob analysis, since the area of the minimum bounding box is too small. In our implementation we require that:

$$A < A_{bb} < B$$

where A_{bb} is the area of the minimum bounding box, A=100, and $B = 4 \times 10^5$ in the real data shown in Fig. F-10.

The blob area analysis threshold keeps isolated dilated clutter regions with a small minimum bounding box from being declared a target, while attempting to mitigate dataframe-wide shifts in the cross-ambiguity that occur during large shifts in the self-ambiguity, as occur even on rock music stations during DJ announcements or a brief quiet period during song transitions. A severe example of such a self-ambiguity shift is depicted in Fig. C2a for an NPR station. One can better appreciate Fig. F-10 as part of a video sequence. In the Results for Blind Target Detection section we present such a short sequence.

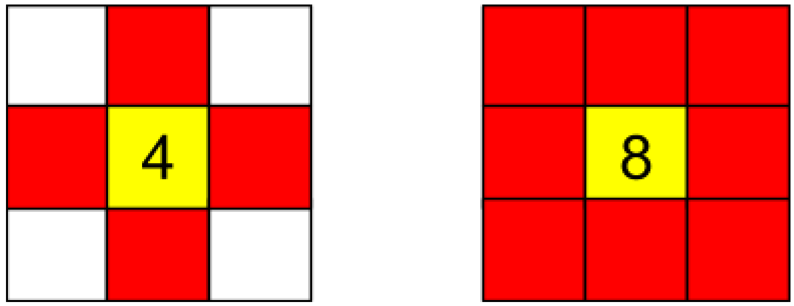

Figure F-8: 4-connected and 8-connected neighbors highlighted in red

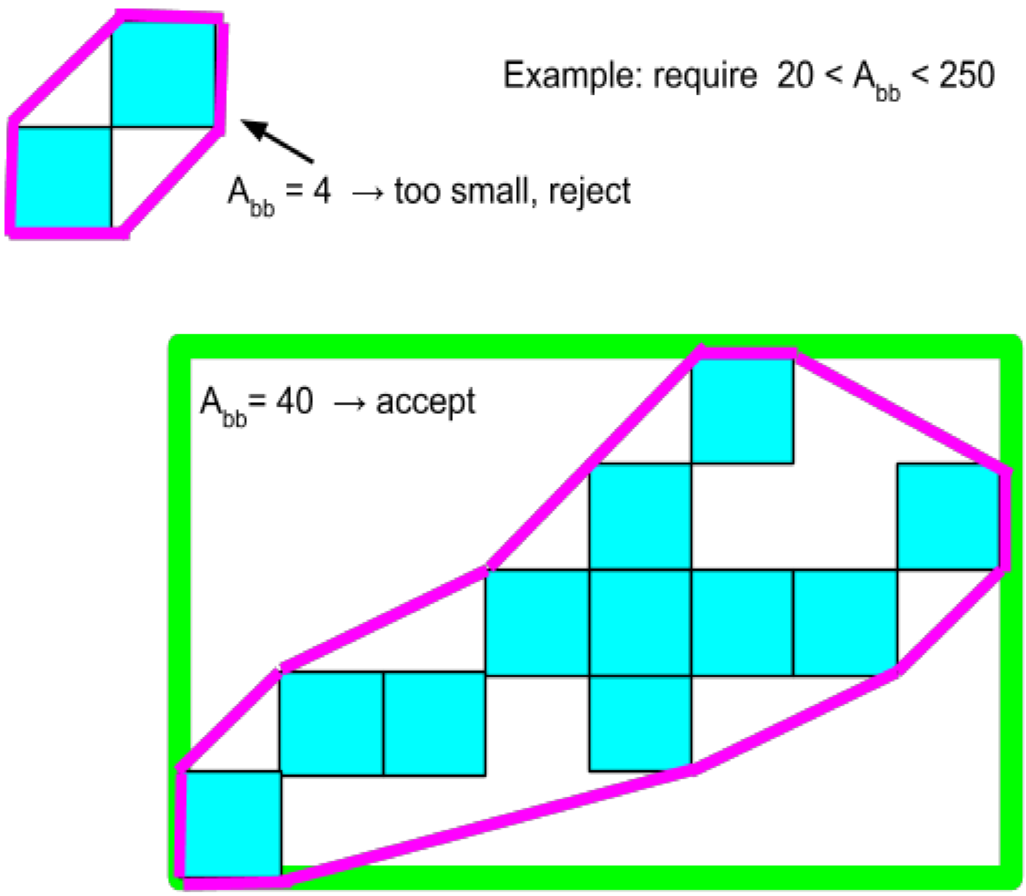

Figure F-9: Cartoon example of connected components and blob analysis

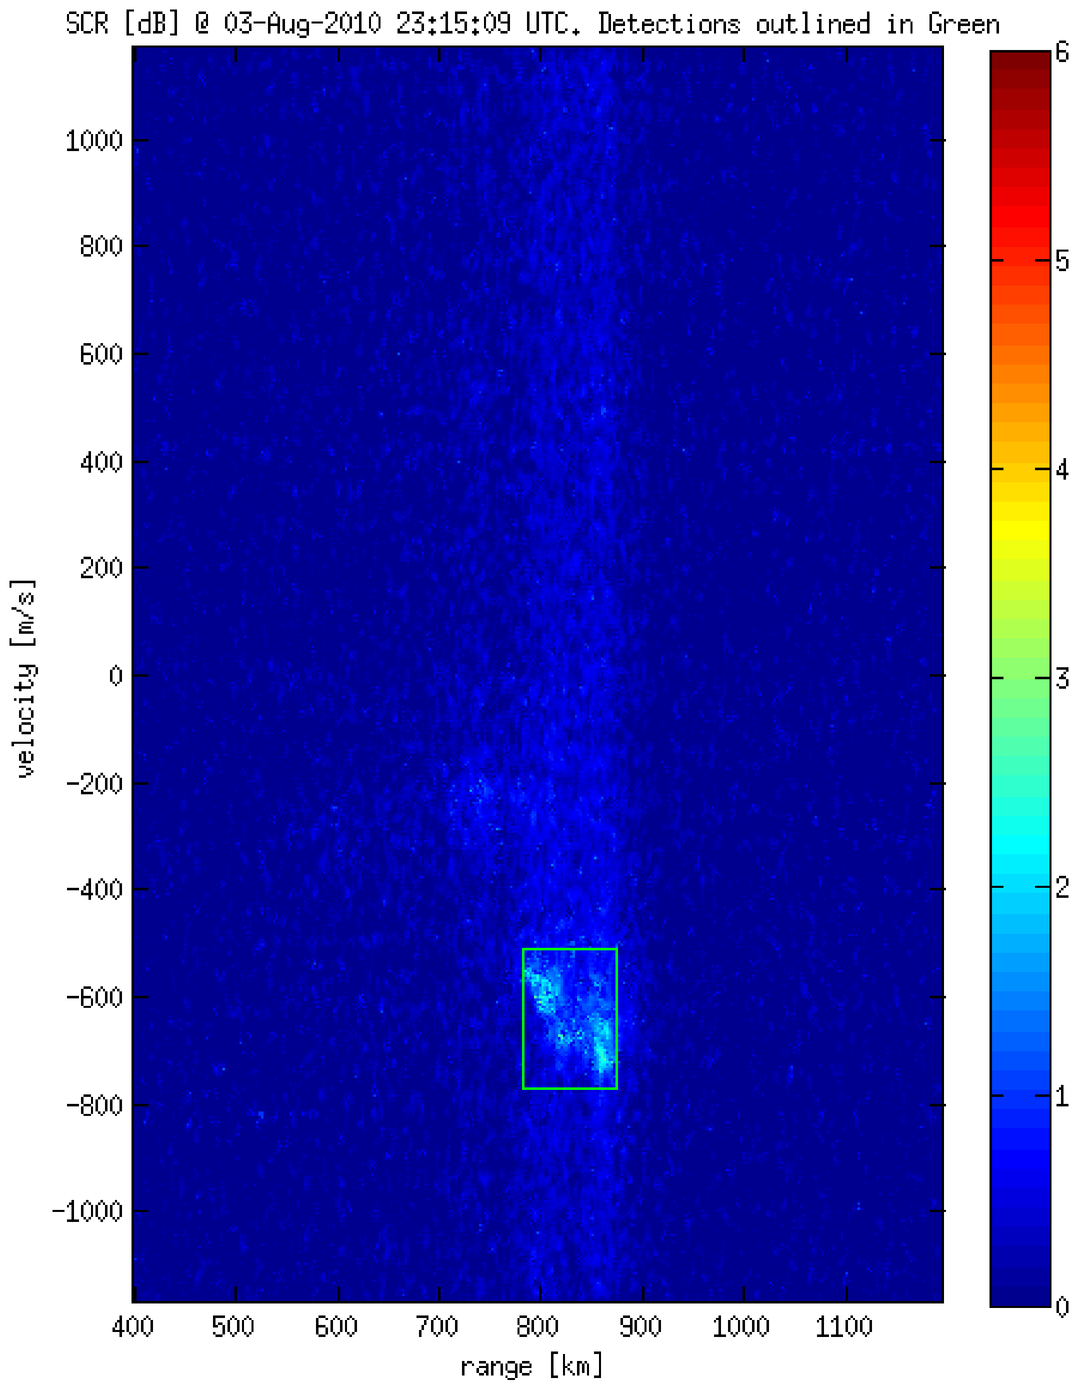

Figure F-10: Example Blob Analysis result on ionospheric returns

RESULTS: BLIND TARGET DETECTION

The cross-ambiguity data used for this first-pass effort were the 2010 August 03 dataset. This data was obtained in the processed state, since the algorithms used as in Fig. D-1 are not publicly available. Figure F-10 shows a correct detection. We do experience rare Type II errors (false negative) as in Fig. G-1, which are a deficiency of this “first-pass” algorithm as we do not yet have memory/tracking as the last step to carry through brief fades in ionospheric clutter returns that cause SCR to drop for a single dataframe. Observe in Fig. G-1b that the erosion operation has aggressively eliminated the weak SCR auroral return, and so the dilation operation will return little if any pixels on the true target return, and so the blob analysis makes a Type II error.

We also experience Type I errors (false positive) as in Fig. G-2, which can occur for the first couple dataframes of a new dataframe sequence as the GMM algorithm has “memory” and has to learn the statistics of a process to create per-pixel Gaussian models best suited to the recent dataframes. Upon major shifts of the self-ambiguity function, the false positives can result from the concomitant large shifts of the cross-ambiguity function on the clutter background. Fig. G-2a shows the GMM output filled with numerous false positives. Many of them are isolated, but enough of them are 8-connected in series to survive the morphological erosion step in Fig. G-2b. This forebodes the ill performance in Fig. G-2c, where the clutter leakage is enlarged into large connected regions by the morphological dilation. The blob analysis will naturally register false positives on the largest associated pixel regions as shown in Fig. G-2d. This only happens for the first frame, and immediately thereafter the memory of the Gaussian fit states is within a range suitable to recognize this clutter as background and no false positives are declared.

Finally, we show a brief sequence of blob analysis in Fig. G-3 showing that sequential frames of the highly dynamic ionospheric returns are detected without false positives. As noted earlier, occasionally there is a single dataframe of false negative, with the possible solutions including Kalman tracking. Observe that for each of the six frames of Fig. G-3, only the aurora is detected. The full video can be viewed at:

http://heaviside.bu.edu/~mhirsch/isis/firstPassDet.avi

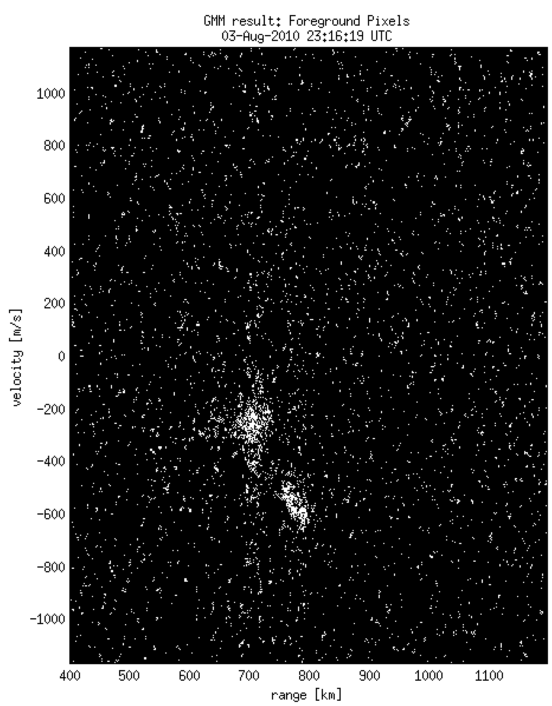

Figure G-1a: GMM output

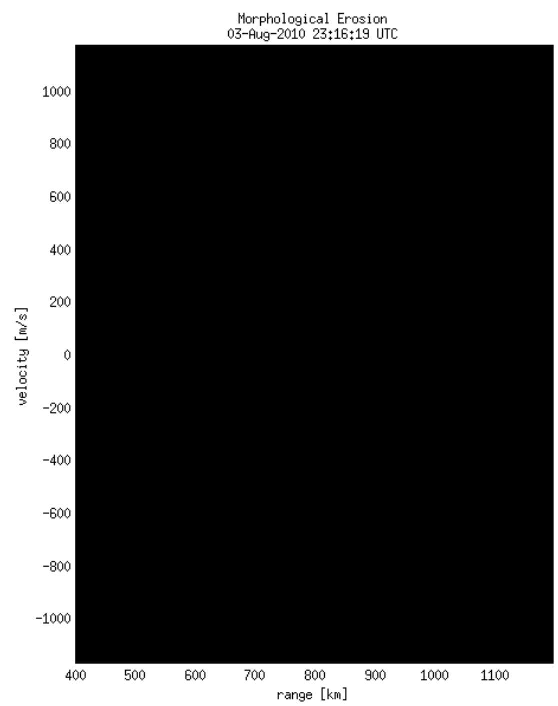

Figure G-1b: Erosion output

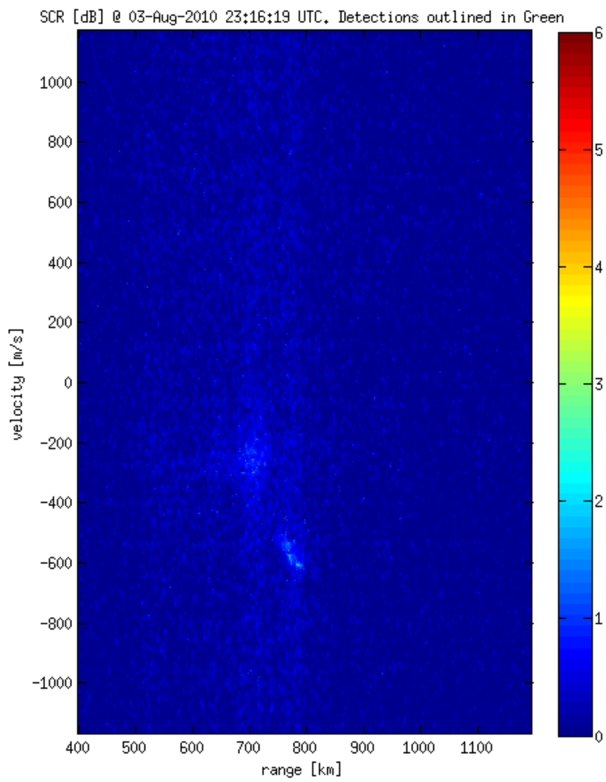

Figure G-1c: Blob Analysis output -- Type II error -- no detections

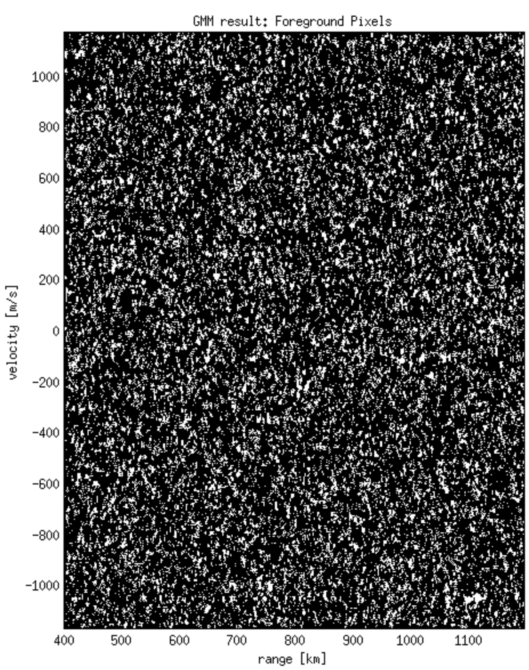

Figure G-2a: GMM output

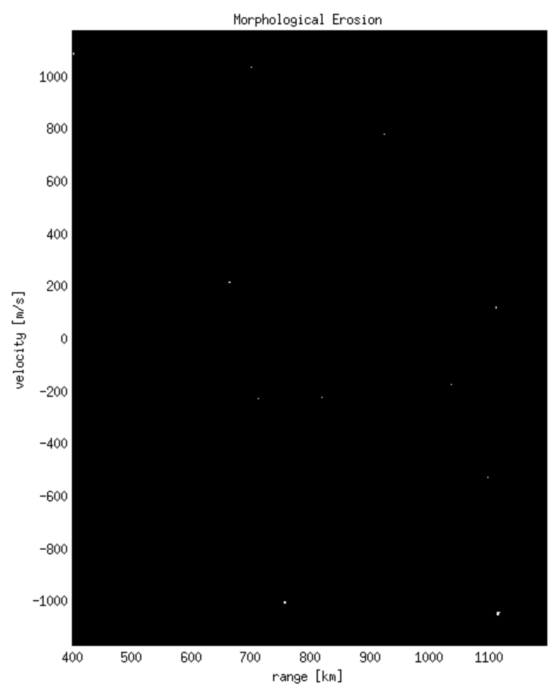

Figure G-2b: Erosion output

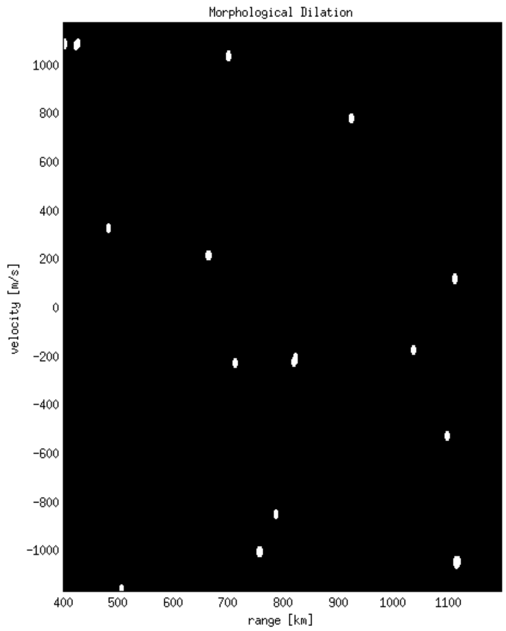

Figure G-2c: Dilation output

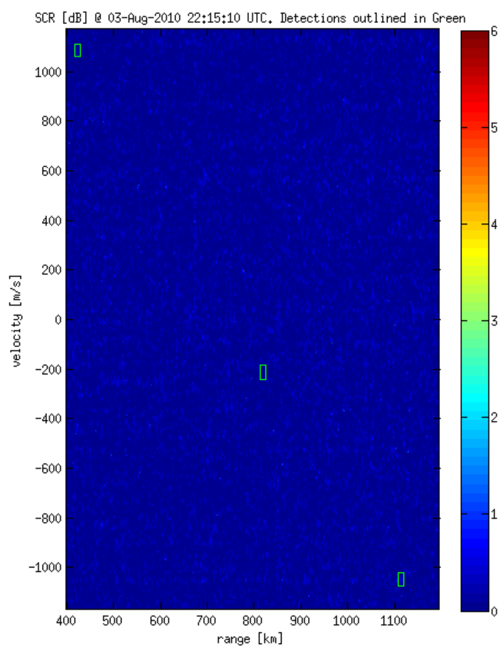

Figure G-2d: Dilation output

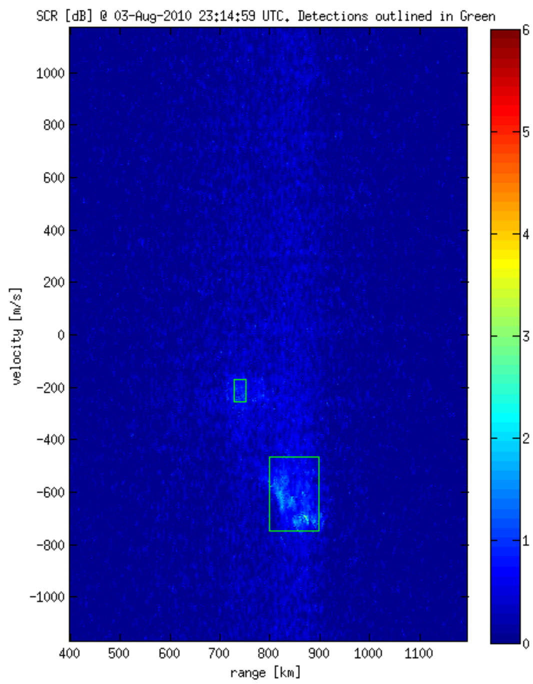

Figure G-3a: Frame 1

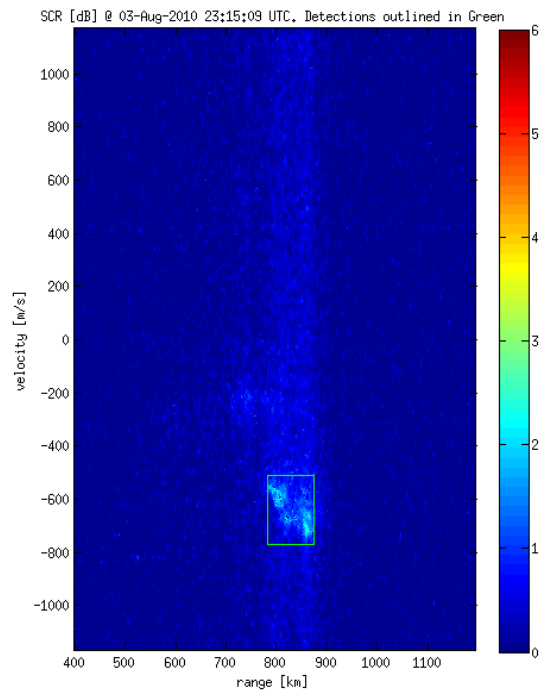

Figure G-3b: Frame 2

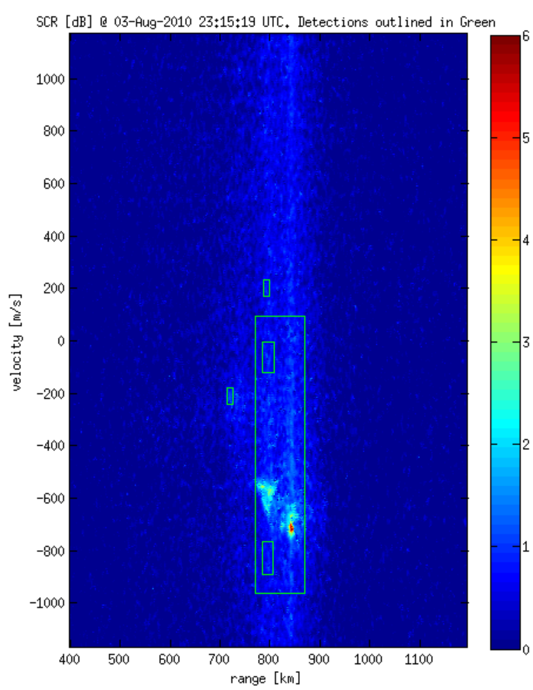

Figure G-3c: Frame 3

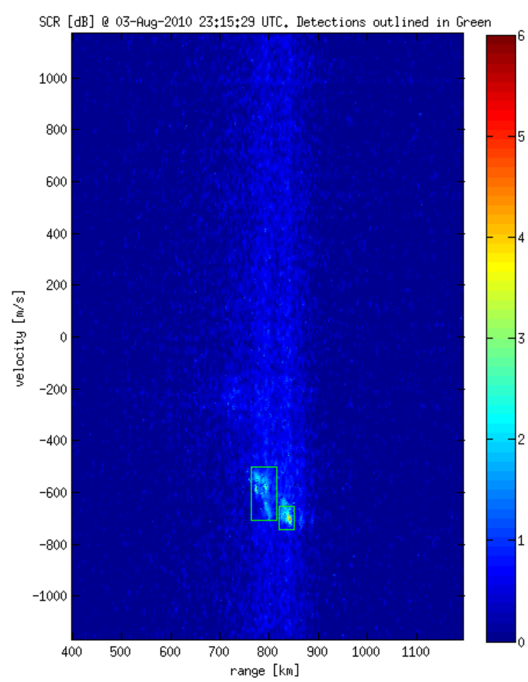

Figure G-3d: Frame 4

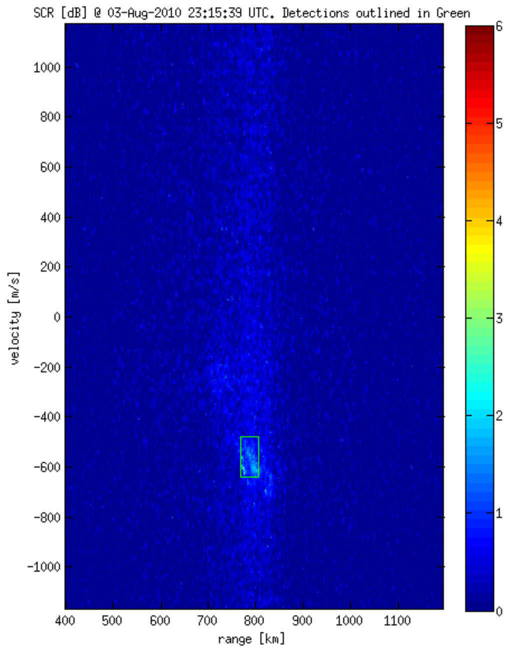

Figure G-3e: Frame 5

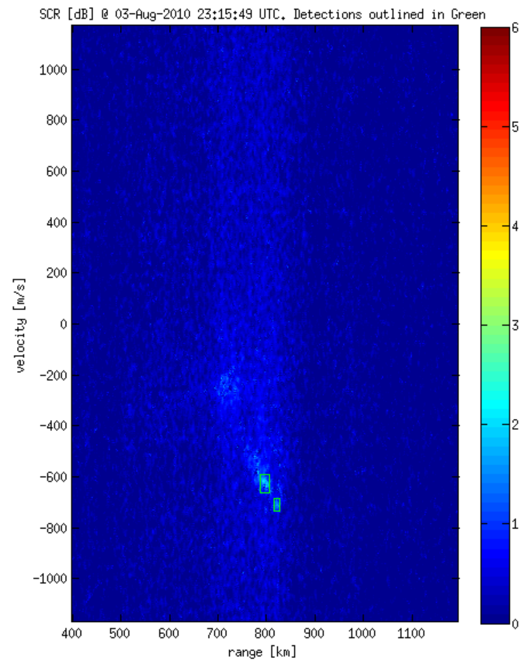

Figure G-3f: Frame 6

The initial results shown in this result have been presented to the project sponsor Frank Lind, and via email Dr. Lind has expressed positive comments about the initial results. The preference would be to have this code in Python as the rest of the ISIS code is written in Python and C. The OpenCV toolbox is available in Python with the same algorithms, and so a future work extension may be to convert the MATLAB code to Python. An obvious future extension would be to verify the algorithm performance on further ionospheric videos, when the ISIS data becomes available for download. Part of the impetus behind better quantifying the self-ambiguity function mentioned earlier in this report was to use this information as qualification on some of the less ideal datasets from NPR and talk radio stations. The data we presented for the machine vision portion of this report was from a rock station, with the relatively stable mean cross-ambiguity shown in Fig. G-4a. An unstable cross-ambiguity is shown in Fig. G-4b, which exhibits discontinuities we would expect from a talk/NPR station. A possible extension is to know which dataframes to discard based on quantified poor self-ambiguity periods.

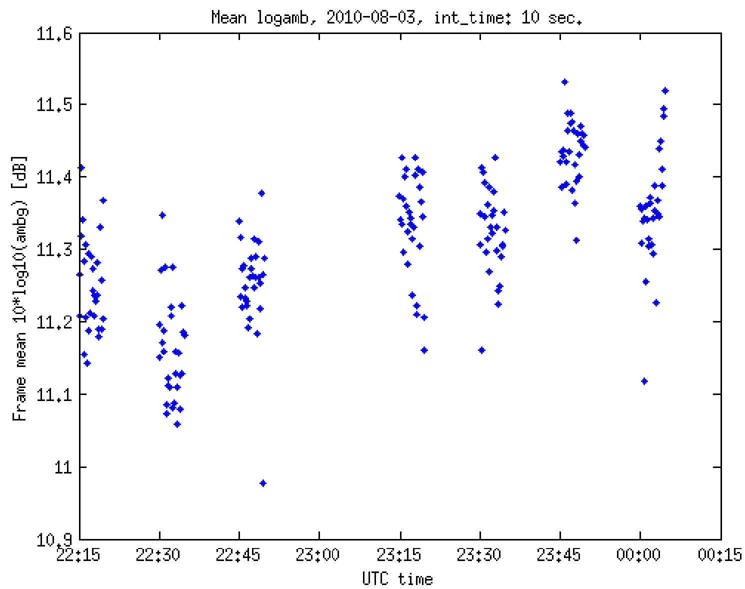

Figure G-4a: stable cross-ambiguity

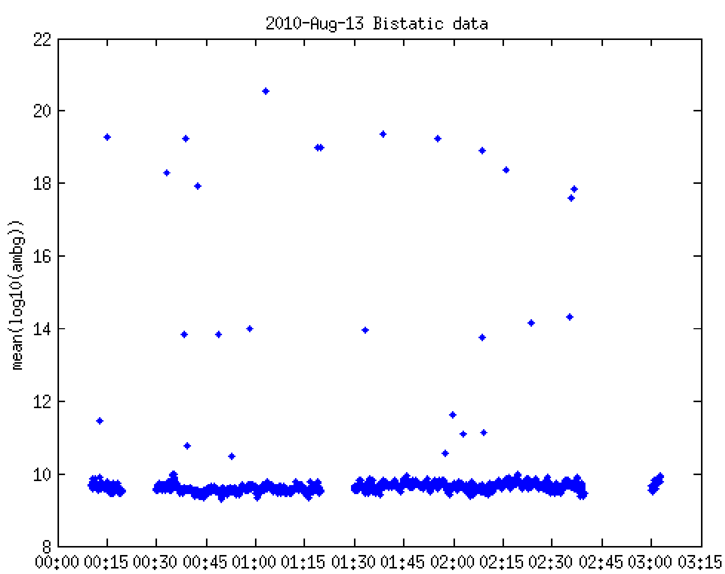

Figure G-4b: unstable cross-ambiguity

CONCLUSIONS: BLIND TARGET DETECTION

We have shown that automated detection of ionospheric returns with high probability of detection is possible with a very low false detection rate. We have not had the opportunity to work with more data yet due to constraints in obtaining data from ISIS. Such limitations may be mitigated by running on one of the main ISIS servers itself--which is necessary in general as much of the passive radar code is proprietary. Without having to extensively tweak parameters -- just with a few observations of the typical range-Doppler extent of a typical ionospheric target, we were able to reliably detect the ionosphere with no false positives beyond the first frame of training (as is expected and can be programmatically discarded). We feel that more data adjacent to times of poor self-ambiguity can be utilized by extending the work in this report. Despite the seemingly endless supply of FM broadcast transmitter/receiver geometries, the high level of interference in the band, only increasing with the transition to HD Radio dictates a need to use stations with less ideal self-ambiguity when using automatic machine vision target detection techniques. We have laid down the first steps in the direction of real-time blind target detection and the customer has seen the initial results favorably.

REFERENCES:

Akbari, H., et al, “Anomalous ISR echoes preceding auroral breakup: Evidence for strong Langmuir turbulence”. *Geophys. Res. Lett.,* vol. 39, is. 3, 2012. doi: 10.1029/2011GL050288

Lind F., “ISIS Annotated Data Examples”, http://atlas.haystack.mit.edu/isis/fermi/events/annotated_data/isis_annotated_data_examples.pdf Aug. 2010, Retreived Nov. 7, 2013.

Lind F., et al, “Intercepted Signals for Ionospheric Science”, *Radio Science*, vol. 48, pp. 248-264, June 2013.

Stauffer, C. and Grimson, W., "Adaptive background mixture models for real-time tracking," *Computer Vision and Pattern Recognition, 1999. IEEE Computer Society Conference on.* , vol.2, pp. 246-252 Vol. 2, 1999. doi: 10.1109/CVPR.1999.784637

Sahr J. and Lind F., “The Manastash Ridge Radar: A passive bistatic radar for upper atmospheric radio science”, *Radio Science*, vol. 32, no. 6, pp. 2345-2358, Nov. 1997.

Sahr J. “Ionospheric Measurements”, Chap. 7 in, *Advances in Bistatic Radar*, SciTech Pub., 2007.

Barnes, A., “The calculation of instantaneous frequency and instantaneous bandwidth”, *Geophysics*, 57, 1520.

Mahafza B., *Radar Systems Analysis and Design using MATLAB*, CRC Press, 2000.

Willis N. and Griffiths H. (eds.), *Advances in Bistatic Radar*, SciTech Pub., 2007.

Cherniakov M. (ed.), *Bistatic Radar: Emerging Technology*, Wiley & Sons, 2008.

Richards M., *Fundamentals of Radar Signal Processing*, McGraw-Hill, 2005.

Gonzalez R. and Woods R., *Digital Image Processing*, 3rd Ed., Pearson Prentice Hall, 2008.

Gonzalez R. et al, *Digital Image Processing Using MATLAB*, 2nd Ed., Gatesmark Pub., 2009.

Meyer M., “Passive VHF Radar Interferometer Implementation, Observations, and Analysis”, M.S. thesis, Elect. Eng. Dept., Univ. of Washington, Seattle, WA, 2003.

Lathi B. and Ding Z., *Modern Digital and Analog Communication Systems.* 4th Ed. Oxford Univ. Press, 2009.

Lee T., *The Design of CMOS Radio-Frequency Integrated Circuits,* 2nd Ed. Cambridge Univ. Press, 2004.

Lyons R., *Understanding Digital Signal Processing.* 3rd Ed., Prentice Hall, 2010.

Carson, J., "Notes on the Theory of Modulation," *Proc. IRE* , vol.10, no.1, pp.57,64, Feb. 1922. doi: 10.1109/JRPROC.1922.219793

Morabito, A., Meyer, M, and Sahr, J.D., "Improved computational performance for distributed passive radar processing through channelised data," *Radar, Sonar and Navigation, IEEE Proc.* , vol.152, no.3, pp.179,184, June 2005. doi: 10.1049/ip-rsn:20045025

Woodward, P., *Probability and information theory*, *with applications to radar,* 2nd ed., Oxford, New York, Pergamon Press, 1963.

Tsao, T., et al, "Ambiguity function for a bistatic radar," *Aerospace and Electronic Systems, IEEE Trans.*, vol.33, no.3, pp.1041-1051, July 1997. doi: 10.1109/7.599331

Howland, P., "Target tracking using television-based bistatic radar," *Radar, Sonar and Navigation, IEEE Proc.* , vol.146, no.3, pp.166-174, Jun 1999. doi: 10.1049/ip-rsn:19990322

Thidé, B., “Simulated Scattering of Large Amplitude Waves in the Ionosphere: Experimental Results”, *Physica Scripta,* vol.T30, pp. 170-180, 1990. doi:10.1088/0031-8949/1990/T30/023

Strømme, A., “Ionospheric and Magnetospheric studies with incoherent scatter radars”, EISCAT/SRI Int., EISCAT School, Qingdao China, 31 Oct 2006. URL: https://www.eiscat.se/groups/Documentation/CourseMaterials/2006Qingdao/Ionospere_China_Anja.pdf Retrieved 2 Dec 2013.

Canfield R., “Central Difference Gradient”, central_diff.m, Air Force Institute of Technology, WPAFB, OH, Oct. 2001. URL:
http://www.mathworks.com/matlabcentral/fileexchange/12-centraldiff-m
Retrieved 8 Dec 2013.

Appendix A: Authorship

While we each contributed to the overall analysis, the primary report writing division of labor went as follows:

Hirsch	Kidd
Introduction	Introduction
BACKGROUND: FM BROADCAST BAND SIGNALS	DESCRIPTION OF METHODOLOGY, RESULTS, CONCLUSIONS: AMBIGUITY FUNCTION ANALYSIS
BACKGROUND: BISTATIC RADAR	
DESCRIPTION OF METHODOLOGY, RESULTS, CONCLUSIONS: BLIND TARGET DETECTION	

APPENDIX B: MATLAB Code

The MATLAB functions are emailed with this report as a ZIP file. We would be delighted to make them available by email/web to the future interested reader.

APPENDIX C: Converting raw binary data to HDF5 MATLAB file

This process is given with commands executed on a Linux computer, but with minor modifications can be executed on a Mac or Windows computer.

1) check that you have installed a Python 2.7 of your choice (e.g. Spyder)
2) Unzip echotek_rfio.tgz into a directory of your choice (e.g. ~/ISIS/python)
3) you have to install each of four Python modules:

```
1. cd rf_raw_io && sudo python setup.py install && cd ..
2. cd PrecisionTime && sudo python setup.py install && cd ..
3. cd file_access_support && sudo python setup.py install && cd ..
4. cd echotek_raw_io && sudo python setup.py install && cd ..
```

4) `cd echotek_raw_io/tools`
5) edit the ecdr_to_mat.py, moving line 89: data_time = float(xb.startSec) + float(xb.startNanoSec / 1.0E9) up to line 84 (within the try statement).
6) Edit lines 70,71 ref_chip and ref_chan to the numbers you're interested in.

Then, to **convert the .bin files in a directory to .mat files** readable by RawReader.m, do (this example for 2010-Aug-03, rx40, 103.5MHz)
ref_chip=0 ref_chan=1

```
python ecdr_to_mat.py -i ~/ISIS/data/2010-08-03/rx40 -o ~/ISIS/data/2010-08-03/rx40
```